

Technical Report 542

Structuring Decisions: The Role of Structuring Heuristics

Patrick Humphreys, Stuart Wooler, and Lawrence D. Phillips

Brunel Institute of Organisation and Social Science,
Brunel University

BASIC RESEARCH




U. S. Army

Research Institute for the Behavioral and Social Sciences

August 1981

Approved for public release; distribution unlimited.



AD A137675

DTIC FILE COPY

84 02 08 008

U. S. ARMY RESEARCH INSTITUTE FOR THE BEHAVIORAL AND SOCIAL SCIENCES

A Field Operating Agency under the Jurisdiction of the
Deputy Chief of Staff for Personnel

JOSEPH ZEIDNER
Technical Director

L. NEALE COSBY
Colonel, IN
Commander

NOTICES

DISTRIBUTION: Primary distribution of this report has been made by ARI. Please address correspondence concerning distribution of reports to: U.S. Army Research Institute for the Behavioral and Social Sciences, ATTN: PERI-TST, 5001 Eisenhower Avenue, Alexandria, Virginia 22333.

FINAL DISPOSITION: This report may be destroyed when it is no longer needed. Please do not return it to the U.S. Army Research Institute for the Behavioral and Social Sciences.

NOTE: The findings in this report are not to be construed as an official Department of the Army position, unless so designated by other authorized documents.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER Technical Report 542	2. GOVT ACCESSION NO. AD-A137675	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) STRUCTURING DECISIONS: THE ROLE OF STRUCTURING HEURISTICS		5. TYPE OF REPORT & PERIOD COVERED Final
		6. PERFORMING ORG. REPORT NUMBER 80-1
7. AUTHOR(s) Patrick Humphreys Stuart Wooler Lawrence D. Phillips		8. CONTRACT OR GRANT NUMBER(s) DAERO-78-G014
9. PERFORMING ORGANIZATION NAME AND ADDRESS Decision Analysis Unit, Brunel Institute of Organisation and Social Science, Brunel Univer- sity, Uxbridge, Middlesex, England		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 2Q161102B74F
11. CONTROLLING OFFICE NAME AND ADDRESS U.S. Army Research, Development and Standardization Group, UK, Box 65, FPO, NY 09510		12. REPORT DATE August 1981
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) U.S. Army Research Institute for the Behavioral and Social Sciences 5001 Eisenhower Avenue Alexandria, VA 22333		13. NUMBER OF PAGES 75
		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE --
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report) --		
18. SUPPLEMENTARY NOTES --		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Decision Making Decision Aiding Heuristics Computer Aiding		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This paper discusses how heuristics can be incorporated in decision aids that are used for structuring decision problems. First, decision theory in terms of linked subsystems is described; this permits identification of the sites where it is necessary to employ heuristic rather than normative, algorithmic procedures. These sites lie on the interface between decision maker's semantic memory, because no adequate formal specification of the structure of semantic memory exists as yet, and consequently no normative (Continued)		

DD FORM 1 JAN 73 1473

EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

1 SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

Item 20 (Continued)

procedures for accessing information from semantic memory can be specified. Next, production systems are used to describe heuristic devices handling transactions across the interface between the decision theoretic system and the decision maker's semantic memory; it is suggested that these heuristics can be incorporated into decision aids designed to improve the quality of access to the information contained within the decision maker's semantic memory. Candidate heuristics for use in development of such decision aids are identified, and criteria for locating the site of operation of heuristics are reported in the literature. Selection of those heuristics that operate on the interface with semantic memory is developed. These heuristics are identified within a comprehensive table of heuristics cited in the literature.

UNCLASSIFIED

Technical Report 542

Structuring Decisions: The Role of Structuring Heuristics

Patrick Humphreys, Stuart Wooler, and Lawrence D. Phillips

Brunel Institute of Organisation and Social Science,
Brunel University

Submitted by:
Robert M. Sasmor, Director
BASIC RESEARCH

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

Approved by:
Joseph Zeidner
Technical Director

U.S. ARMY RESEARCH INSTITUTE FOR THE BEHAVIORAL AND SOCIAL SCIENCES
5001 Eisenhower Avenue, Alexandria, Virginia 22333

Office, Deputy Chief of Staff for Personnel
Department of the Army

August 1981

Army Project Number
20161102B74F

Basic Research

Approved for public release; distribution unlimited.

ARI Research Reports and Technical Reports are intended for sponsors of R&D tasks and for other research and military agencies. Any findings ready for implementation at the time of publication are presented in the last part of the Brief. Upon completion of a major phase of the task, formal recommendations for official action normally are conveyed to appropriate military agencies by briefing or Disposition Form.

STRUCTURING DECISIONS: THE ROLE OF STRUCTURING HEURISTICS

BRIEF

Requirement:

To investigate if and how subjective judgmental heuristics can be incorporated into a computer interactive decision-making aid.

Procedure:

Review and critiques of the following topics were made: the development of decision-theoretic systems, decision theory, heuristic devices, and classification of heuristics. A computer aid was developed using heuristics to support the structuring and evaluation of a decision or problem.

Findings:

The current view of judgmental heuristics as biases is challenged on the grounds that the biasing nature of some of the heuristics has not been replicated in the degree, and, in some cases, in the direction of the bias. An alternative is to view heuristics as devices for integrating the decision maker's semantic memory with the decision-theoretic techniques. In addition, it was possible to identify areas in the decision-making process in which heuristic rather than normative algorithmic procedures should be used. A comprehensive categorization of heuristics was generated. A decision aid, Multiattribute Utility Decomposition (MAUD), is used to demonstrate how heuristic devices can be used to handle the transactions across the decision-theoretic computerized system and the decision maker's semantic memory. Specifics on the MAUD system are summarized in Technical Report 543.

Utilization of Findings:

This report will be of use to those interested in understanding the psychological means of supporting decision making. Due to their technical and exploratory nature, the materials presented will be of limited value for immediate application. Rather, their merits will be realized as a stimulus for researching both the empirical psychological support and the efficacy of such concepts in decision aiding.

STRUCTURING DECISIONS: THE ROLE OF STRUCTURING HEURISTICS

CONTENTS

	Page
1. SUMMARY	1
2. INTRODUCTION	1
Interfacing Machine-Client-Decision Analyst: What Next? . . .	1
3. THE PROVINCE OF DECISION THEORY	2
Overview	2
3.1 Historical Development of Decision-Theoretic Systems . . .	3
3.2 Consequences of Extending the Province of Decision Theory	11
4. HEURISTIC DEVICES	13
Overview	13
4.1 A Production System Representation	14
4.2 A Critique of Heuristics as Biases	15
4.3 Heuristics as Biases, A Case Study: Conservatism, Misperception, or Misaggregation	18
4.4 Heuristics Within Production Systems: Analysis of Representativeness	20
4.5 Classification of Heuristics	27
5. AN APPLICATION	27
Overview of MAUD	27
5.1 Example of the Use of a Heuristic Device: Use of Structuring Heuristics in the Elicitation of Poles of Attribute Dimensions	27
5.2 Production System Control of Structure Ordering and Reordering Activities	38
6. STRUCTURING UNCERTAINTY	43
6.1 Current Systems	44
6.2 Buffering for Functional Relationships	46
6.3 Modular Uncertainty Structure	47
CONCLUSION	53
REFERENCES	59

CONTENTS (continued)

	Page
APPENDIX A: PRODUCTION SYSTEMS AND SEMANTIC MEMORY	67

LIST OF TABLES

Table 1. Classification of heuristics	28
2. Modular uncertainty structure	49
3. Sites of operations of structuring heuristics	57

LIST OF FIGURES

Figure 1. Inputs and outputs involved in interfacing a decision-theoretic event-act linkage to a "real world" problem . . .	5
2. PIP system inputs and outputs	10
3. Extended province of decision theory: core system buffered by three subsystems	12
4. General form of the setup required for the investigation of bias in human information processing	16
5. Model of the cognitive processes involved in inference . .	22
6. Block diagram of Tversky and Kahneman's (1974) model of judgment biased by representativeness	24
7. Interaction between decision maker and MAUD using the Difference Method to elicit poles of an attribute dimension	37
8. Productions used to control preference structure ordering and reordering activities within MAUD	40
9. Ways of assessing uncertainty about some target event . . .	44
10. Modular uncertainty structure for hull claims model	48
11. Diagram of basic operations required for creating a decision structure adequate for evaluation by the decision-theoretic system	55

1. SUMMARY

This paper discusses how heuristics can be incorporated in decision aids used for structuring decision problems. First, we describe decision theory in terms of linked subsystems; this permits us to identify the sites in which it is necessary to employ heuristic rather than normative algorithmic procedures. These sites lie on the interface between decision-theoretic subsystems and the decision maker's semantic memory, because no adequate formal specification of the structure of semantic memory exists as yet, and consequently no normative procedures for accessing information from semantic memory can be specified. Next, we make use of production systems to describe heuristic devices handling transactions across the interface between the decision-theoretic system and the decision maker's semantic memory and suggest that these heuristics can be incorporated into decision aids designed to improve the quality of access to the information contained within the decision maker's semantic memory.

We identify candidate heuristics for use in development of such decision aids. We develop criteria for locating the operation site of heuristics reported in the literature and then select those heuristics that operate on the interface with semantic memory. These heuristics are identified within a comprehensive table of heuristics cited in the literature.

2. INTRODUCTION

Interfacing Machine-Client-Decision Analyst: What Next?

A fundamental problem in decision analysis is the extent to which the decision maker should work directly in interaction with a decision analyst, or instead, interact with an automated decision aid in building, exploring, and evaluating a decision-theoretic model of the problem being considered. Should decision aids be interfaced directly with the decision maker, or should the decision analyst serve as an intermediary, retaining control of the overall structuring of the problem, but using a decision aid to perform the more "automatic" functions (such as applying a multiattribute utility composition rule, rolling back a decision tree, and so on)? Initially, automated aids were designed with the goal of "bootstrapping" the decision maker (Dawes & Corrigan, 1974). Automating a decision maker's composition rule provided judgments superior to his or her unaided intuitive judgments. Subsequently, it was realized that structuring and decomposition operations might also be profitably aided (Selvidge, 1976). Humphreys (1978b) found that it was the structuring capability of a bootstrapping aid that was valued, particularly by those decision makers who were initially unsure about their preferences and goals.

However, designing a decision aid with structuring capability does not involve simply an extension of the devices employed in aids designed to automate composition rules. As we show in section 3, decision theory specifies completely how the elements of a decision are to be combined, but it provides only a limited set of rules about how the elements should be structured (e.g., events should be mutually exclusive and exhaustive). Thus, heuristic devices

must be made available to guide structuring activities. These heuristics must operate within a comprehensive control and checking framework to insure that the final structure is valid and internally consistent.

The research program reported here was concerned with both the design of a structuring aid and the identification of heuristics that can facilitate structuring activity. In addition, we considered new approaches to structure problems involving uncertainty. Specifically, the original proposal that led to this work called for the following tasks to be carried out:

1. Develop interactive Multiattribute Utility Theory (MAUT) aid with dynamic capability,
2. Collect and investigate heuristics useful for predecisional structuring, and
3. Identify and program useful probability assessment procedures.

Soon after we started this work, a new conceptualization of decision theory began to emerge. This new view helped us to see where decision aids would be most useful, and it provided criteria for judging which heuristic would facilitate structuring activity and which would not; this view will be discussed in section 3. Then, in section 4, we will discuss the role of heuristics and present a list of heuristics, derived from a wide-ranging examination of the literature, along with evaluation of their usefulness for structuring. Section 5 presents Multiattribute Utility Decomposition (MAUD), the interactive structuring aid whose development was completed under this project. Section 6 reports on work in developing new approaches to uncertainty structuring. In section 7 we will conclude with an overview of the structuring process and suggest uses for heuristic devices within this process.

3. THE PROVINCE OF DECISION THEORY

Overview

Our view of decision theory is that it consists of four interlocking systems. One is concerned with the representation of utility, another with the representation of uncertainty, a third (the core system) with modeling act-event linkages, and the last with the influence of secondary events not explicitly modeled in the core system. Decision theory provides for the coherent interfacing of these systems, so that an output from one system across an interface completely specifies the value of the input to the system on the other side of the interface. For example, the value of a holistic utility of a consequence as calculated using MAUT can be used in expected utility calculations when folding back an act-event tree.

Although decision theory prescribes how these systems should link together, it is still necessary to provide inputs from outside; in this sense, decision analysis is not a closed system. This is always the case, whether decision analysis is viewed as an engineering science or as a clinical art (Buede, 1979). Models of the environment provide inputs in the engineering science approach,

whereas judgments by experts are the main inputs in clinical art applications. Subjective judgments necessarily involve semantic memory systems (Anderson, 1976). However, the province of decision theory (Lindley, 1974) does not include models of semantic memory nor of the environment.

In the remainder of section 3 we will outline briefly the historical development of decision theory into the four interlocking systems. In section 4, we will demonstrate how procedures that involve passing information across an interface between a system modeled within decision theory and the decision maker's cognitive information processing system require the use of rules specifiable only as heuristic devices, outside the axiomatic formulation of decision theory.

3.1 Historical Development of Decision-Theoretic Systems

Decision theory has, since its inception, been concerned with the decomposition of immediate acts (Bernoulli, 1737).¹ Acts can be decomposed into the set of their possible consequences, each described by a payoff function indexing the value gained or lost through the realization of that consequence. Given the assumption that the decision maker will choose a specific subsequent act in each eventuality under consideration, the decision analysis can be conducted by constructing a payoff matrix. Each cell in this matrix defines a particular act-event sequence, the consequence of which is represented by a payoff shown in the cell. This is known as the normal form of decomposition of immediate acts through act-event sequences.

An alternative decomposition of immediate acts through act-event sequences is known as the extensive form, usually represented as a decision tree. Here the links between all immediate acts, events, subsequent acts, and consequences under consideration are linked in a tree structure containing act nodes and event nodes. The left-hand side of the tree starts with a simple immediate act node, and the right-hand side ends with a set of consequences. The set of payoffs associated with these consequences can be used to fill the cells of a payoff matrix identical to that obtained through a normal form decomposition of the same problem. The difference, however, is that an extensive form structuring of the problem allows one to fold back the decision tree piecemeal, examining expected utilities at intermediate nodes in the tree, and providing expected utilities for immediate acts.

Hence, we can see that normal form and extensive form² decompositions of a set of immediate acts into events, subsequent acts, and consequences differ only in that the extensive form provides the decomposition in a more structured way, allowing an analysis of the problem in partially decomposed form, as well as in holistic and fully decomposed form. Both the normal and extensive forms

¹ Bernoulli was concerned with, among other things, the question why people ever answered the question, "Should I buy insurance now?" in the affirmative.

² See Luce and Raiffa (1957, chapter 3) for a detailed description of normal and extensive form decomposition.

are what we will call single-system decompositions, and in this section we shall trace the history of decision theory, starting with these single-system decompositions, up to its present state of development in which decision problems can be subjected to decompositions by four separate but interlocked systems.

At the outset of this development, about 20 years ago, initial research emphasis was on the nature of the optimal composition rule for use within the single decomposition system in prescribing a preference function over the set of immediate acts, so that the optimum immediate act could be chosen under any circumstances (barring hindsight). Later, as a consensus began to develop around the advocacy of the subjective expected utility (SEU) rules as the normatively preferred composition rule (Edwards, Lindman, & Phillips, 1965; Lee, 1970), the emphasis shifted to the nature of the optimal structuring of the decision problem (e.g., Brown, 1977) in conjunction with the question, How should the structure adopted be extended or constrained in the light of difficulties likely to be encountered in interfacing it with the decision maker and the environment?

Figure 1 shows the interfaces involved, together with the inputs and outputs crossing these interfaces.

Within personalist decision theory, the decision-theoretic model is not interfaced directly with the environment, but environmental inputs and outputs are mediated through the decision maker's judgments. This process involves the invocation of semantic memory in the provision of "knowledge of the world" (Anderson & Bower, 1974). This point has often been missed, sometimes resulting in the construction of decision analytic models proposed for interfacing directly with the environment (Howard, 1966), despite the realization that subjective judgments are involved right down the line (Staël von Holstein, 1977).

Decision theory does not in itself provide any model of semantic inputs, or of the environment, so in each case inputs cross an interface from a system not formally modeled to one that is. Hence, the total system cannot be closed, and the decision-theoretic model must be conceived as part of an open system in which elicitation instructions are sent as outputs across the interface requesting the necessary information to proceed with modeling the decision problem.

Three different types of information are required for input with the decision-theoretic model:

1. Information about influences of acts on events, and vice versa;
2. Information about probabilities to be assigned at event nodes; and
3. Information about utilities of consequences.

Because the system on the other side of the interface is not formally modeled, no axiomatic way of specifying the form of the output (elicitation instruction) will guarantee that the input subsequently received is that required at that point in the generation of the decision-theoretic system. Hence any input

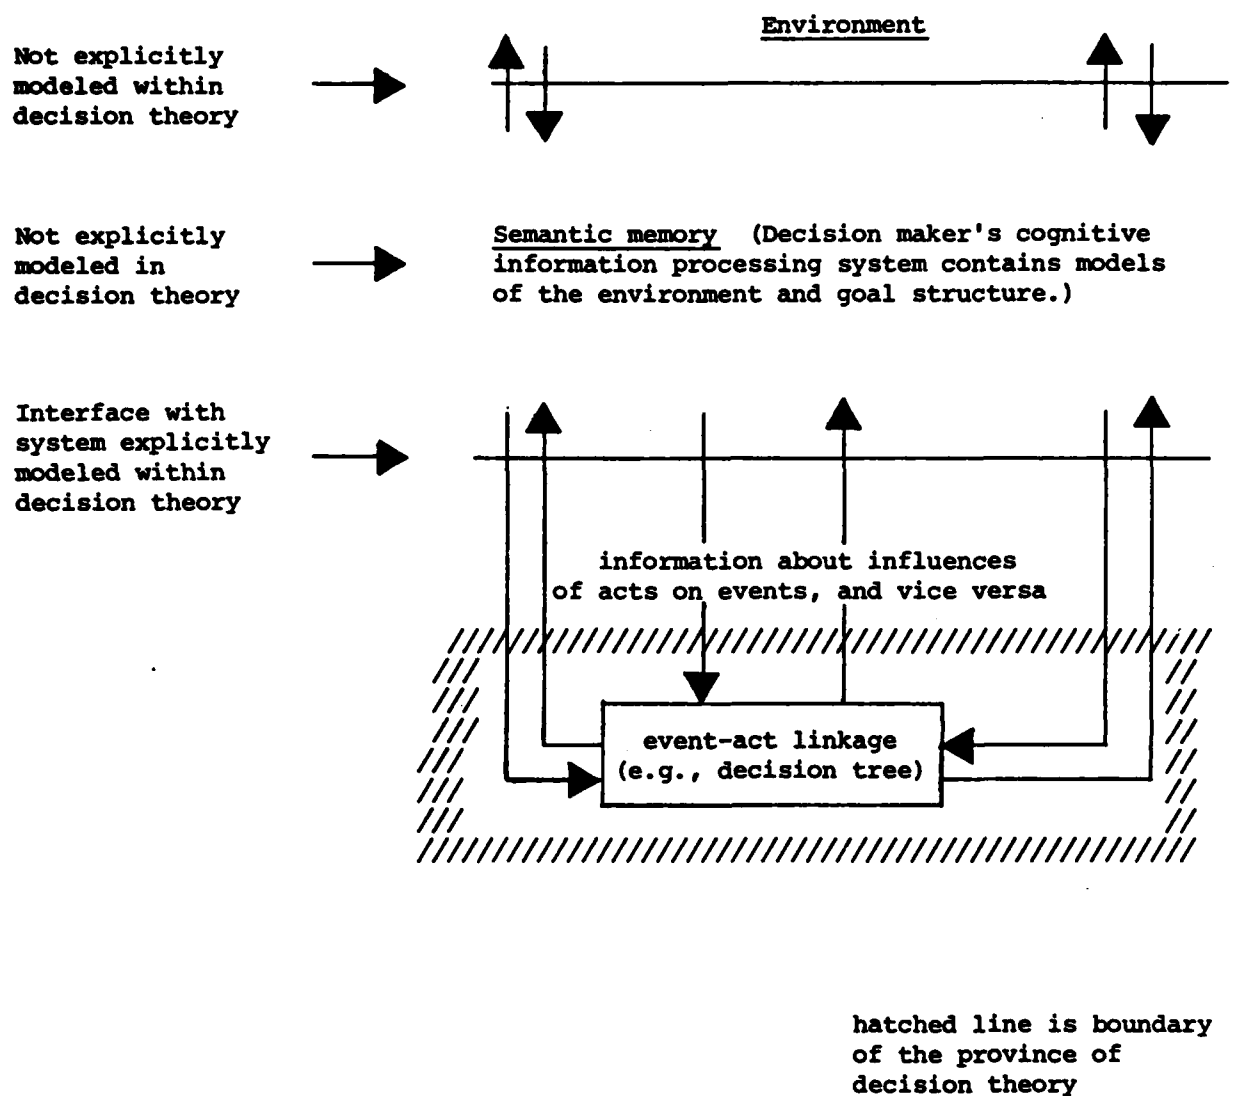


Figure 1. Inputs and outputs involved in interfacing a decision theoretic event-act linkage to a "real world" problem.

must, if possible, be checked for coherence before continuing with its use in modeling the problem. The lack of formal criteria for identifying valid inputs into decision-theoretic models explains why decision analysts have been uncharacteristically vague in specifying elicitation instructions to be used in obtaining any particular input. Decision analysts have resorted to exhortations such as, "the decision analysts should take on a role not too dissimilar to that of a psychoanalyst" (Brown, Kahr, & Peterson, 1974).

Despite the fact that optimal elicitation instructions cannot be specified axiomatically, some forms of elicitation instruction will be more efficient than others. In subsequent sections of this report we will develop a production system representation of heuristics designed to optimize such efficiency. First, though, we will describe revisions and extensions of the decision-theoretic system shown in Figure 1 resulting from research enterprises during the past two decades, that were, in effect, designed to push back the interface between act-event linkage system and the decision maker's semantic memory through buffering the system with three further systems.

In the original representation outlined in Figure 1, input information about act-event links is specified in terms of binary (presence or absence) relationships; input information about probabilities is specified in terms of (unidimensional) scalar numbers assignable to consequences in a one-to-one relationship. Eliciting these various inputs in applied problems has not always been easy, so decision analysts have developed systems, specified axiomatically within decision theory, designed to buffer these inputs. The general specification of each of these three buffer systems is considered below.

Buffering Act-Event Linkages: The Influence Diagram

The basic decision-theoretic model of Figure 1 assumes that the decision maker will choose a specific subsequent act in each eventuality under consideration. This assumption implies a binary linkage between events and subsequent acts: A specific act is either consequent on the immediately preceding event in the decision tree structure, or it is not. However, in attempting to elicit such linkages, decision analysts often received the nonbinary input, "it depends." Selvidge (1975, p. 46) gives an example:

Suppose someone were to ask you whether or not you will buy a new car next year. You might answer that it depends on the state of repair of your old car. At the present you do not know what the outcome of this secondary³ event "state of repair of your old car" will be

What can the decision analyst do about this? Selvidge continues:

... but you can list the different outcomes and then for each of these assess the probabilities that you would or would not buy a new car under those circumstances.

³ For "secondary," read "not explicitly modeled within the act-event tree."

Selvidge then shows how such information can be structured through the use of an influence diagram, modeling an act (or event) whose occurrence is conditioned by a number of secondary events and defines the composition rule appropriate for use within the influence diagram structure.⁴

The core system describes linkages between acts and events leading to consequences. Influence diagrams, in effect, act as a buffer for this core system, enabling the modeling of secondary events that influence the assessments to be made within the system.

Buffering Utility Assessment: Multiattribute Utility Theory

When the value of a consequence can be completely described in terms of money (e.g., worth \$1,000), utility assessment may be reasonably conceived as mapping a unidimensional utility function on to monetary value. The scalar monetary value is then transformed into a scalar utility value, which is input to the basic decision-theoretic model. However, many consequences possess moderately or extremely complex value structures (e.g., "have a child," Beach et al., 1976). Attempts to elicit scalar holistic utility assessments of such consequences directly from decision makers are usually unsuccessful. The decision maker typically responds to the elicitation instruction with the reply that "there are too many factors to trade off." Early attempts to solve this problem were based on elaborating the act-event tree into the future, looking for subsequent consequences that would possess a simpler value structure and that would therefore facilitate direct scalar assessment of their utilities. However, there is no guarantee that such an elaboration will uncover consequences with simpler value structure, and the elaboration has the additional undesirable effect of pushing the decision horizon further into the future, a future that exists only as a fantasy in the decision maker's mind and that may, at the time an immediate act has to be considered, not be modelable with any precision (Brown, 1978; Humphreys, 1979).

An alternative, and usually more efficient, solution is to stay with consequences possessing complex value structures in the act-event tree but to buffer the inputs representing utility assessments of these consequences using MAUT to provide a further decomposition of their value structures (Raiffa, 1969; von Winterfeldt & Fischer, 1975; Humphreys, 1977).

The structure of this decomposition may be modeled in either normal form or extensive form. In normal form decompositions, the structure of each consequence is decomposed into part worths (Kneppreth, Gustafson, Leifer, & Johnson, 1974) on a number of attribute dimensions. For example, the consequence of building a particular type of rapid transit system may be decomposed into travel time, user comfort, vehicle construction cost, user fatalities, level of environmental noise, etc. (Raiffa, 1969). The input to the MAUT system is vectors of part worth (decomposed utility) assessments from the decision maker (or expert), and a MAUT-axiomatized composition rule is applied to these

⁴ For a specification of computer-assisted procedures for use in such modeling, see Allen et al., 1976.

vectors to yield the holistic utility values for consequences (Humphreys, 1977). These scalar values are the output of the MAUT system and the input to the act-event system with which it is interfaced.

In extensive form decompositions, a utility hierarchy is constructed, in which the holistic utility of each consequence is decomposed within a tree structure (here called a hierarchy). At the bottom level of the hierarchy the decomposed part worth assessments are input, the same inputs required by a normal form decomposition. However, the composition of these inputs into holistic, scalar utilities of consequences (emerging at the top of the hierarchy) is performed in stages, by multiplying through the hierarchy (equivalent to folding back a decision tree), permitting the examination of partially decomposed utilities of consequences in addition to the fully decomposed input assessments and holistic utilities. Examples of the use and interpretation of such hierarchical representations of the utility structure of consequences are given in Beach, Townes, Campbell, and Keating (1976); Chinnis, Kelly, Minkler, and O'Connor (1975); and Fischer, Edwards, and Kelly (1978). In addition, the interactive multiattribute utility decomposition and recomposition decision aid, MAUD, computes and uses a hierarchical utility structure to aid in eliciting weights on the various attribute dimensions.

All of these systems buffer the act-event structure. They do this by moving the site of the utility interface with semantic memory, so that the relevant elicitation instructions are designed to elicit inputs representing decomposed utilities of consequences on attribute dimensions, rather than holistic utilities.

Buffering Probability Assessment: PIP (Probabilistic Information Processing) Systems

The information about probabilities to be assigned to event nodes, which is required as input to the act-event system shown in Figure 1, is of the form of the probability of a hypothesis given particular data: $P(H|D)$.

The hypothesis (H) is that the future event represented at a particular event node in the decision tree will occur. The data (D) summarize the information, not shown on the tree, relevant to the event and given the intervening scenario represented by the linkage within the decision tree connecting that event to the present situation (immediate act). Edwards (1962) and Pitz (1975), among many others, have pointed out that estimates of the required $P(H|D)$ are often not readily available and have to be constructed by integrating $P(D|H)$ over a variety of data and hypotheses. For example, consider the hypothesized event, "enemy launches attack." One must estimate $P(H|D)$: the probability that this event will happen, given D, the state of affairs at that future point in time at which the decision maker believes the hypothesized event may occur. Faced with the need to estimate this, the decision maker usually starts considering data defined in terms of $P(D|H)$: What is the probability associated with particular states of the world given that the enemy actually launches an attack? With a complex world, a large number of $P(D|H)$ s exist; hence a wide range of data may have to be considered and integrated in the attempt to obtain a reasonably well-defined assessment of $P(H|D)$.

Pitz (1975) outlines methods by which a person may do this within the structure of his or her own semantic memory. For an appropriate model of semantic memory, Pitz used that of Anderson and Bower (1973), a precursor of systems developed by Anderson (1976), which are discussed in Appendix A. For our purposes we need only to note that such procedures almost invariably produce suboptimal results, often grossly so, when compared with those readily obtainable through the use of Bayes' theorem to perform the integration, within what Edwards (1962) called a Probabilistic Information Processing (PIP) system. Edwards' definition of a PIP system is outlined in Figure 2.

The input to the PIP system (elicitation instruction) across the interface from the decision-making system is a request for information about the likelihood of an event ($H?$) and data about the states of the world obtaining at the relevant (future) time, D . The output returned across this interface is $P(H|D)$, which becomes an input to the act-event system. In order to obtain $P(H|D)$, the PIP system requires estimates of $P(D|H)$ over all H s and D s relevant in establishing $P(D|H)$. These are elicited either from decision makers (or experts) or from formal (computer-based) models of the environment by asking them to consider the probability of the data if H were to obtain. Memory is involved, because the data that need to be considered are almost certainly not those impinging at the present moment on the decision maker's senses, or the machinery implementing the formal model's sensors, but have to be recalled from earlier experience, yielding data believed to be relevant in establishing the state of the world, D , under consideration.

Hence, use of a PIP system serves to buffer inputs concerning uncertainty about future states of the world to the act-event system. PIP systems were originally proposed as labor-saving devices for cases where $P(H|D)$ was not directly estimatable, providing Bayes' theorem as an automated composition rule. However, the superiority of Bayes' theorem over intuitive composition rules for integrating $P(D|H)$ s to obtain $P(H|D)$ was soon confirmed in a wide variety of contexts (Edwards & Phillips, 1964; Edwards, Phillips, Hays, & Goodman, 1968; Howell, 1967; Gustafson, 1969), and led to the suggestion that such buffering should be included in decision-aiding systems where possible (Slovic & Lichtenstein, 1971). The Bayesian composition rules for use in PIP systems were originally formulated for normal form operation, that is, where the composition of the input information into $P(H|D)$ is accomplished in a single stage. Extensive form representations were developed later, for use where the composition is accomplished through intermediate stages within a hierarchical tree structure (Kelly & Barclay, 1973). This technique, known as cascaded inference, has been found useful in situations in which it is possible to decompose the problem of estimating $P(H|D)$ through the use of intermediate or explanatory variables, i.e., where

... it will often be possible to assess the likelihood of the data given some intermediate variable, and the likelihood of that intermediate variable given another, and so on, until the hypotheses of interest are reached. (Kelly & Barclay, 1973, p. 388)

Schum and Kelly (1973) have also developed cascaded inference composition schemes for determining the inferential impact of confusing and conflicting reports from a mixture of unreliable sources, where these reports provide data diagnostic with respect to a particular $P(H|D)$ under consideration.

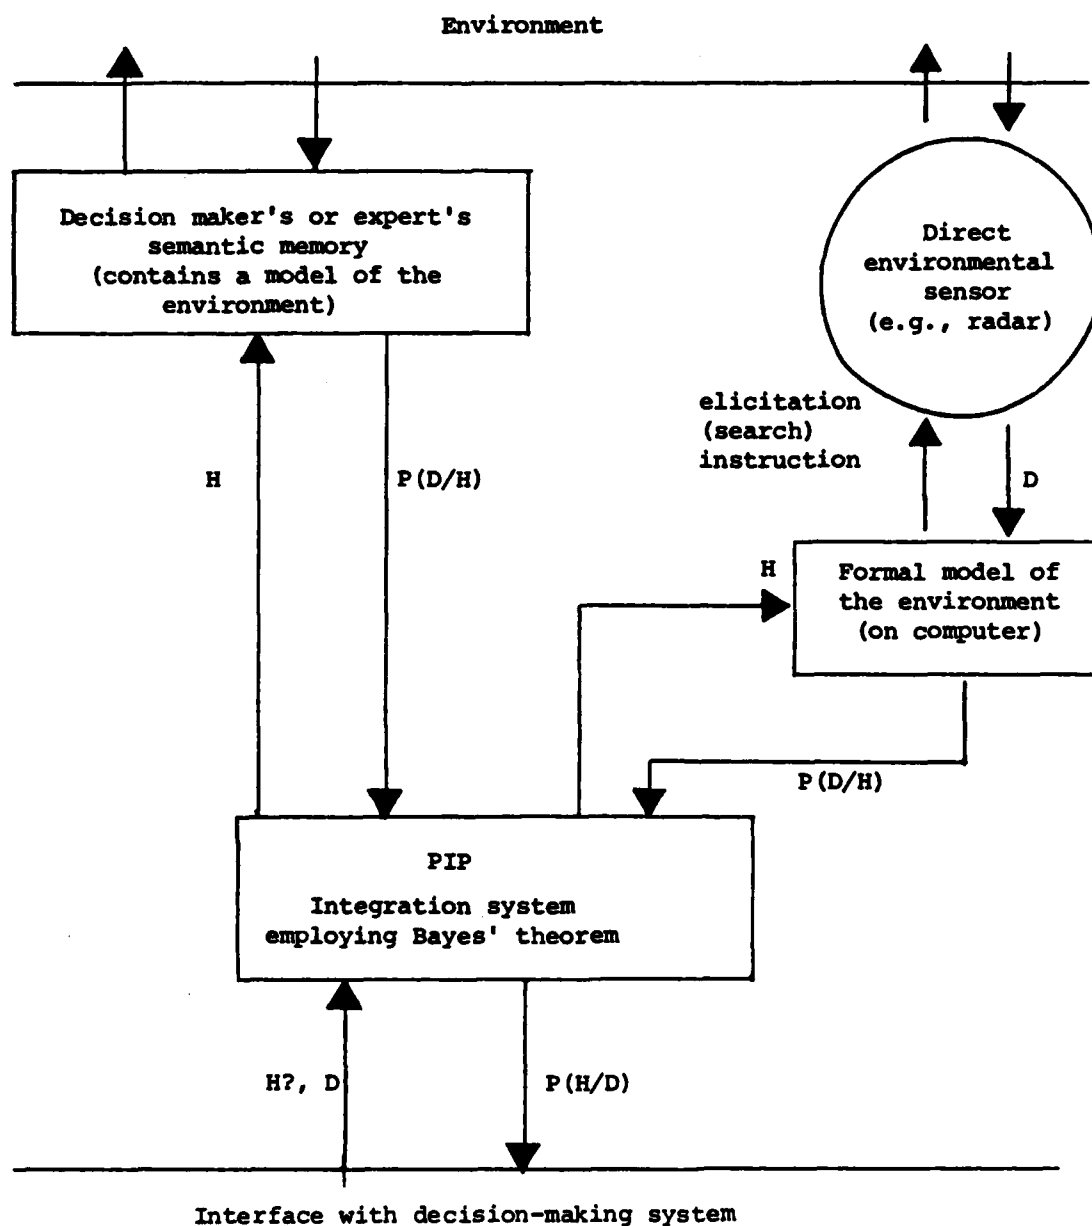


Figure 2. PIP system inputs and outputs.

From the point of view of this report, all such systems, whether expressed in extensive or normal form, extend the province of decision theory in the same way: by providing a subsystem to buffer the inputs concerning probabilities to be assigned to events within the structure of the core act-event system.

3.2 Consequences of Extending the Province of Decision Theory

The developments described in section 3.1 serve effectively to extend the province of decision theory by buffering all three interfaces of the core act-event system of decision theory through subsystems involving the decomposition of the relevant inputs. The inputs of the subsystem are still interfaced with the decision maker's semantic memory, just as were originally the inputs of the stand-alone act-event system.⁵ However, the decomposed information now obtained through these inputs (part-worth assessments, $P(D|H)$ assessments, and influence relations between primary and secondary events) are integrated into the relevant utility, probability, and event inputs to the core system by composition rules applied within the subsystems. Figure 3 summarizes these relationships.

The province of decision theory has thus been extended to include four subsystems within the total system bounded by the hatched line shown in Figure 3. Each subsystem is constructed on the basis of decision-theoretic axioms and is (or should be) coherent and explicitly specified in a way consistent with these axioms. The nature of transactions between each subsystem and the core act-event system is also completely explicit, because the effect on the relevant input to the core system of modification of content or structure within any buffer subsystem is completely and exhaustively specifiable on the basis of the relevant decision-theoretic axioms. However, this specificity is not true for transactions that cross the interface between any of the three subsystems and semantic memory, because only one side of the interface is explicitly modeled within decision theory. We are not able to say precisely what effect a particular modification of content or structure within semantic memory will have on the resulting input from memory to a particular decision-theoretic subsystem (or vice versa), because any precise specification of a person's semantic memory structure is, of necessity, missing.⁶ Nevertheless, while there is no axiomatic way of specifying transactions across the interface

⁵ The PIP systems were also proposed for direct interfacing to formal models of the environment, as shown in Figure 2. Here we are concerned with decision-making situations in which no such fully structured formal model of the environment is available a priori, and in which the structuring task facing the decision analyst involves what Humphreys (1979) called internal ordering and reordering (as opposed to external ordering through the use of a formal model). In such cases the interface of the decision-making system is always with the decision maker's (or expert's) semantic memory.

⁶ Formal models of semantic memory (e.g., Anderson & Bower, 1973; Kintsch, 1974; Norman, Rumelhart, & CNR Research Group, 1975; Anderson, 1976) provide precise specifications of systems which are incomplete and fragmented approximations to semantic memory structure, intended for investigative purposes only.

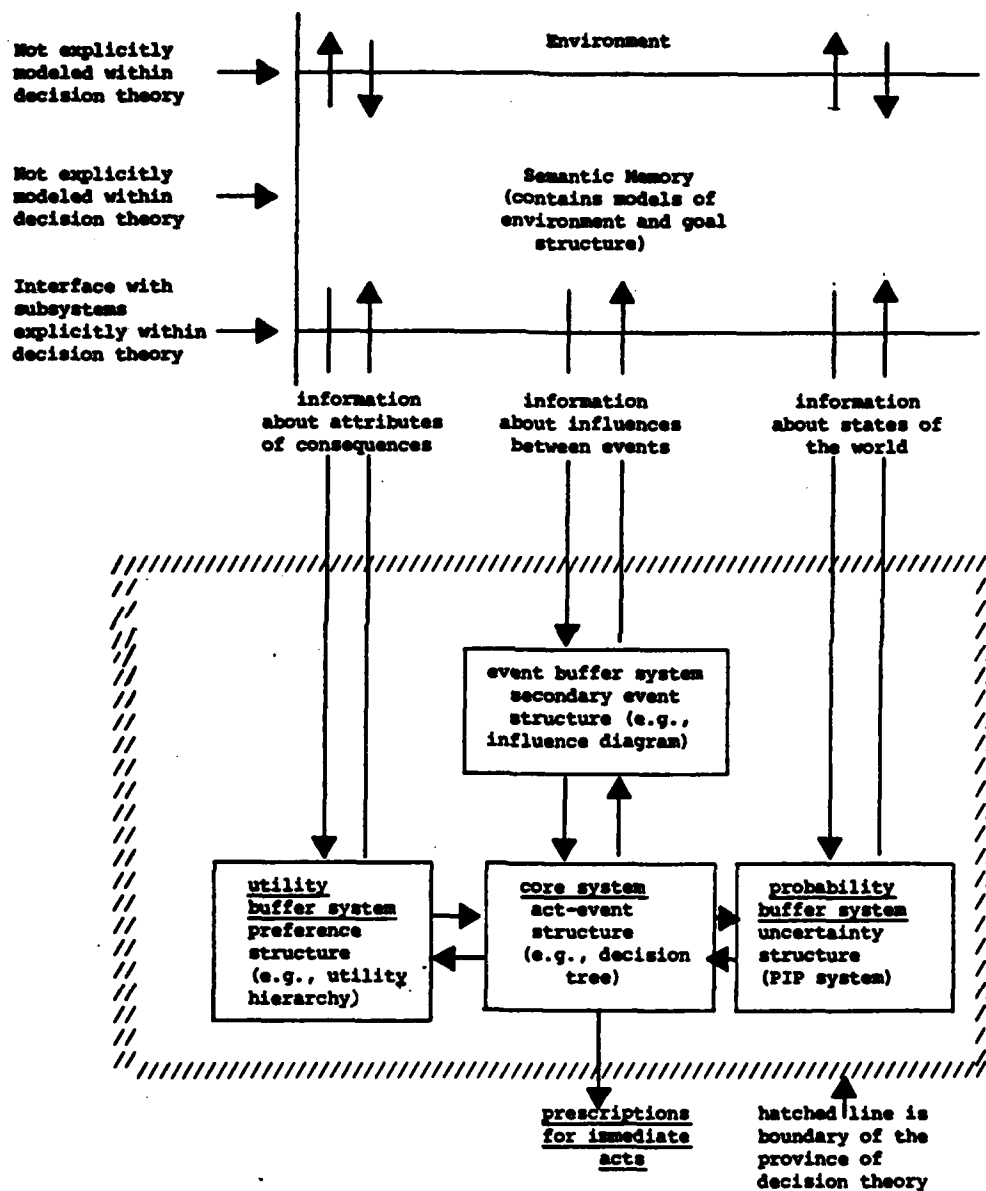


Figure 3. Extended province of decision theory:
core system buffered by three subsystems.

with semantic memory, this paper advances the thesis in subsequent sections that heuristic devices have a major role to play in optimizing such transactions to meet the demands of the decision-theoretic system.

4. HEURISTIC DEVICES

Overview

Heuristic devices by definition are suboptimal when compared with normative formulations for the same operations. Working within a normative framework, the only justification for the use of heuristics (and resulting risk of inefficient use of information or incoherence of the material generated) is that they are "quick and easy" or "reduce information processing load" and hence generate results that would not otherwise be available because the decision maker would refuse to invest the amount of time or effort involved in following a normatively acceptable procedure. In this report, we are not interested in pursuing such a line of justification;⁷ rather, we are concerned with conditions in which it is not possible to specify a procedure on a normative basis, leaving no option but to resort to a heuristic device. This occurs only in situations in which the procedure involves crossing an interface between a system that is explicitly modeled and one that is not. The only interfaces that meet this specification in Figure 3 are the three between the decision theory subsystems and memory.

For example, one of these interfaces is crossed when a set of criteria is elicited from the decision maker for purposes of assessing the utilities of consequences within the utility buffering system shown in Figure 3. Given that it is not possible to specify the procedure to be used here in normative terms, what happens in practice? According to Humphreys (1979),

Practitioners forced to think about the problem have tended to draw on analogies with problems facing psychoanalysts and clinical psychologists. Brown, Kahr and Peterson (1974) suggested that the decision analyst take on a role "not too dissimilar to that of a psychoanalyst," although such a suggestion, taken seriously might involve more than they had bargained for (c.f., Sandler, Dare & Holder, 1973). Keeney (1975) described a MAU decomposition used in studying alternative energy policies by presenting the verbatim record of a session with the decision maker in case report format. Humphreys and Humphreys (1975) suggested the use of elicitation techniques designed for use within repertory grid technique, which was originally developed at Harvard psychological clinic (Kelly, 1955). None of these techniques, of course, are grounded in any axiomatic theory of preferences. They are all able to elicit structural material that would not have been volunteered without their use, and it

⁷ For an example of a comprehensive attempt of such justification of particular heuristic procedures enclosed within a decision-theoretic framework, see Wallsten (1978).

is possible to check coherence of the structures so generated. However, there is no guarantee that the resulting coherent structures are in any way optimal.

Humphreys (1979) discusses the need for internal reorderings of preference structures generated through the use of heuristic techniques to bring them closer to optimality while maintaining coherence. Section 5 describes the structuring (and restructuring) capability of the computer program MAUD, whose operation is controlled through the operation of a production system. Section 4.1 provides an outline introduction to the notion of a production system and the potential uses of such systems within decision-theoretic contexts.

4.1 A Production System Representation

Since Emile Post's pioneering work (Post, 1943) on a powerful, new symbol manipulation system--which he called a production system--several authors have developed Post's idea as a basis for the specification of psychological models of human knowledge.⁸

One of these authors, Newell (1973), explains the basic operations of production systems as follows:

A production system is a scheme for specifying an information processing system. It consists of a set of productions, each production consisting of a condition and an action.... A production system, starting with an initially given set of data structures, operates as follows. That production whose condition is true of the current data (assume there is only one) is executed, that is, the action is taken. The result is to modify the current data structures. This leads in the next instant to another (possibly the same) production being executed, leading to still further modification. So it goes, action after action being taken to carry out an entire program of processing, each evoked by its condition becoming true of the momentarily current collection of data structures.

The advantage of production system representations of information processing systems is their great generality and flexibility. As Newell and Simon (1972, p. 835-6) report:

Methods are to be represented as production systems, and the set of all methods is equivalent to the set of all production systems that will realise rational courses of action for some given goals and some environment. A specific problem solver has available some repertory of such methods, which come to control his behaviour under various conditions.

⁸ For a general review of the use of production systems, see Davis and King (1975).

A number of authors have taken up Newell and Simon's suggestion of the application of production systems to problem-solving tasks and have generated models of the productions implemented within a person's internal judgmental or semantic memory system. The best articulated model is that of Anderson (1976), which is reviewed in Appendix A. However, the only application within a decision-theoretic context is that of Pitz (1979), who describes how production systems can be used to generate inferences made by subjects through testing conditions of representativeness (Tversky & Kahneman, 1974).

The general form of a production rule⁹ is

$$R : \{c \rightarrow a\}$$

where c is a list of conditions that may be met, and a is a list of actions to be taken consequent on c being met. In generating a production system, Pitz was primarily interested in providing an account of how his subjects might use such rules in ordering and noting the contents of their semantic memories to arrive at inferences. However, in this paper we are primarily interested in how to develop such systems for use in directing and controlling decision aids in interaction with decision makers.

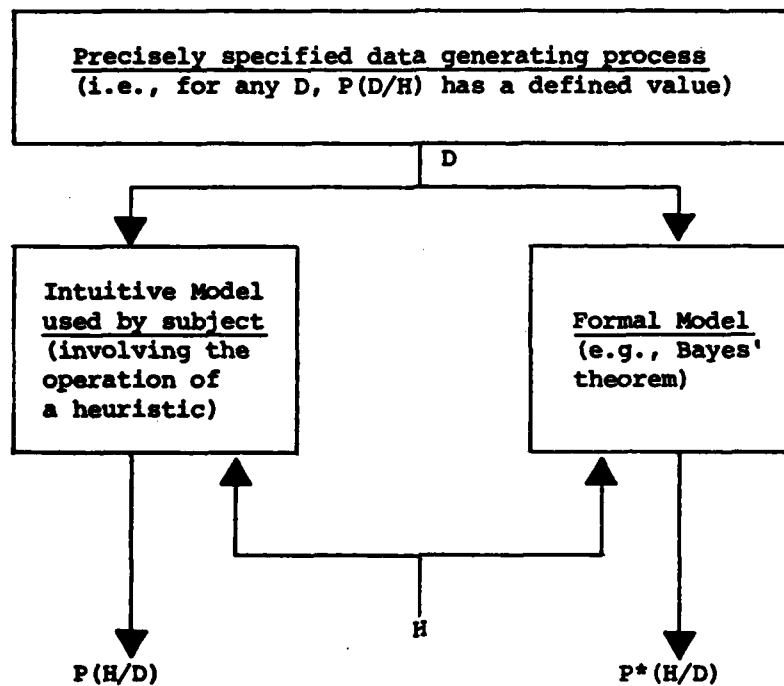
What would such decision aids be like? They would, at the very least, possess the following features:

1. They would give elicitation instructions to activate (but not control) the decision maker's semantic memory and to elicit specific outputs;
2. They would check for coherence within the material thus output across the interface to a decision-theoretic subsystem (explicitly modeled in the decision aid) and take whatever action is necessary should such checks fail.

4.2 A Critique of Heuristics as Biases

The study of the use of heuristics within decision-making contexts has traditionally been linked with the study of bias in human information processing. Attribution of bias is possible only in situations in which it is possible to compute what an unbiased response (inference) might be. The requirements for such situations are outlined in Figure 4. The degree of bias is computed by comparing the output of the intuitive model being used by the subject (person suspected of being biased) with the output of a formal model axiomatically grounded within decision theory (e.g., Bayes' theorem) when both are supplied with identical input data. The larger the difference, the greater

⁹ In section 4 we use the notation $P : \{C \rightarrow A\}$ in preference to $R : \{c \rightarrow a\}$ to show (i) that the Rule (R) is a production rule (P) and that C and A are vectors.



$P^*(H/D)$ is taken as the criterion value, and the difference between $P(H/D)$ and $P^*(H/D)$ forms the basis for the measure of bias.

Figure 4. General form of the setup required for the investigation of bias in human information processing.

the bias. In an attempt to insure that the input to the two models is identical, the data-generating process has to be precisely specified.¹⁰

The intuitive model employed by the subject is presumed to comprise a number of heuristic procedures, and the degree of suboptimality of such heuristics is indexed by the degree of computed bias, that is, the degree of deviation from the optimal results prescribed by the output of the formal model.

Two aspects common to treatments of heuristics as biases are irrelevant to the issues addressed by this paper. These are

1. The construction of a precisely specified data-generating process. We are interested in situations in which much of the data has to be retrieved from the decision maker's semantic memory (see earlier section), and in which there is no precise external specification of the structure of the processes involved in the generation of such data.
2. The comparison with a formal model axiomatically based in decision theory. In the specification and implementation of decision aids, we assume that where such models exist, and where the intuitive models used by (unaided) decision makers yield output that is sub-optimal compared with that provided by the formal model, then the decision aid will implement the formal model in preference to the intuitive model,¹¹ and consequently, the resulting judgment will be "unbiased."

As described in section 2, what we are interested in is the use of heuristics at sites in which no formal model can be axiomatically specified, and in which, therefore, the issue of suboptimality cannot be investigated by a direct test against a formal criterion. We will have to address the question of optimality among heuristics, not optimality of heuristics.

However, the first step in such an investigation is to assemble a set of heuristics that can serve as candidates in our investigation of problem structuring and that may therefore merit incorporation within decision aids. The criterion for including a heuristic in this set is that it must be sited at the interface between a decision-theoretic subsystem and semantic memory (see Figure 3).

¹⁰ A typical specification is as follows: "Chips drawn at random, with replacement from two bags, each containing a large number of chips, but with a defined composition of chips, such as 70% red and 30% blue, or 30% red and 70% blue in each bag"; see Phillips and Edwards (1966).

¹¹ This rationale for decision aids is called "bootstrapping" (Dawes & Corrigan, 1974); see also Humphreys, 1977 (section 6.1), and Humphreys, 1979 (section 3).

In section 4.3 we present a case study of the investigation of two processes purported to underlie the reported conservatism bias in probabilistic information processing,¹² misaggregation and misperception. The case study details the isolation of the heuristics involved in these accounts of bias and demonstrates how they may be rejected from our set of candidate heuristics because they are not sited at the required interface.

In the literature, consistent deviations from the output of an appropriate formal model are often attributed to a global causal mechanism such as "availability" or "representativeness" (Tversky & Kahneman, 1974). In such cases, we found that a reformulation of the postulated processes using a production system representation, like that discussed in section 4.1, was a necessary step in isolating the actual heuristic procedures involved prior to examining their site of operation.

In section 4.4 we present a reformulation of Tversky and Kahneman's account of representativeness in these terms.

Section 4.5 summarizes the results of our formulations of the existing literature and provides a classification of heuristics along the following lines:

1. Effects of heuristics upon intuitive judgments, as reported in the literature;
2. Explanation suggested by the authors reporting the effect; and
3. A statement of which of these explanations of heuristic effects we intend to investigate further for use within decision-aiding systems employing structuring heuristics, and our reasons for doing so.

4.3 Heuristics as Biases, A Case Study: Conservatism, Misperception, or Misaggregation?

Conservatism in intuitive probabilistic judgment was described by Slovic and Lichtenstein (1971) as follows:

Upon receipt of new information, subjects revise their posterior probability estimates in the same direction as the optimal model but the revision is typically too small: subjects act as if the data are less diagnostic than they truly are.

These conservative responses are commonly accounted for by one of two explanations: as the result of either intuitive misperception of the data generator (Lichtenstein & Feeney, 1968; Peterson, DuCharme, & Edwards, 1968; Pitz & Downing, 1967; Wheeler & Beach, 1968; Vlek & Beintema, 1967) or misaggregation of various pieces of information to produce a single holistic response (DuCharme & Peterson, 1967; Peterson & Swensson, 1968; Hammond, Kelley, Schneider, & Vancini, 1967).

¹² For research establishing this bias, see, for example, Peterson, Schneider, and Miller (1965), Phillips and Edwards (1966), and Pitz and Downing (1967).

The misperception thesis is explained by Slovic and Lichtenstein (1971) in the following way:

In order to perform optimally subjects must have some understanding of the data generator, the model, device, equation or assumptions used by the experimenter to generate the stimuli shown to the subject. If the subject misunderstands the data generator he may misperceive the conditional probability of the data given the hypothesis, $P(D|H)$.

In the laboratory experiments commonly used to support the misperception thesis, the true nature of the data generator (such as a particular type of probability distribution) is specifiable precisely because the experimenter generates the stimuli shown to the subject from well-specified sources such as that described in footnote 10. In real world contexts, it is difficult to specify the true nature of the data generator, because in such real world cases it is notoriously difficult, if not impossible, to specify a veridical model of the task environment with which intuitive models can be contrasted.¹³ Misperception becomes a misnomer when veridical perception cannot be specified, but the explanation may still be of interest because it indicates a feature of methods used to model stimuli within the individual's information processing apparatus.

But where is this feature located within the system shown in Figure 1? It is neither a feature of the assessment of content within decision structure (although it may lead to biased judgments) nor of a structuring heuristic used to access information from semantic memory for input into the decision structure. Rather, it is a feature of subjective encoding of information given in the environment within the network of semantic memory. It thus operates at the interface between the environment and semantic memory, and not at the interface between semantic memory and decision-theoretic subsystems, which is of concern to us here. Hence we can reject misperception as a structuring heuristic useful within our terms, because it is not situated at the appropriate interface in Figure 2.

On the other hand, the misaggregation thesis of information processing interprets conservatism as the result of intuitive inadequacy in information aggregation when compared to the procedures for aggregation prescribed by the Bayesian method. On this explanation, therefore, conservatism is the consequence of a feature of intuitive assessment of information within the probability buffering subsystem shown in Figure 3 and is thus of no help to us in specifying ways in which subjects structure the decision space itself, because once again it is not situated at the appropriate interface.

In summary, while both the notions of misaggregation and misperception lead to the specification of heuristics, these heuristics did not meet the criteria for inclusion in the set of heuristics that might be useful in decision structuring operations, because they were not sited at the appropriate interface. The next section, however, describes the isolation of a heuristic

¹³ See quote from Newell and Simon in section 4.4 to understand why this is so.

which, although conceived within an account of bias in information processing, is located at the appropriate interface and has already been reported elsewhere in use as a decision-structuring heuristic (Sheppard, 1976; Humphreys, 1979, section 5.2).

4.4 Heuristics Within Production Systems: Analysis of Representativeness

This section presents in some detail a production system representation of one heuristic, described by Tversky and Kahneman (1974), which we believe is a component in intuitive structuring of certain kinds of decision problems. Tversky and Kahneman named it the "representativeness" heuristic. We delve fairly deeply into the modeling of this example to give a detailed demonstration of a production system representation of a structuring heuristic. The subsequent sections of this paper (and our work emanating from it) will be centered on the implications of structuring heuristics, rather than their detailed modeling. First, though, we detail the modeling to clarify the function of structuring heuristics and to show the advantages of their specification as production systems.

The effects accredited to the use of the representativeness heuristic have been shown to bias intuitive judgment considerably and to be highly replicable. Tversky and Kahneman (1974, p. 1124) have introduced it as follows:

Many of the probabilistic questions with which people are concerned belong to one of the following types: What is the probability that object A belongs to class B? What is the probability that event A originates from class B? What is the probability that process B will generate event A?

In answering such questions, people typically rely on the representativeness heuristic, in which probabilities are evaluated by the degree to which A is representative of B, that is by the degree to which A resembles B. For example, when A is highly representative of B, the probability that A originates from B is judged to be high. On the other hand if A is not similar to B, the probability that A originates from B is low.

A consequence of evaluation by representativeness is that "if people evaluate probability by representativeness, therefore, prior probabilities will be neglected" (Tversky & Kahneman, 1974, p. 1124). Tversky and Kahneman have tested this effect in a variety of experiments, one of which is examined here in detail. In this experiment subjects were asked to judge the likelihood that an individual, call him Dick, is either a lawyer or an engineer. Subjects were given two sets of data from which to make their judgment:

1. Base-rate data: lawyers and engineers are in a ratio of 7:3 in the population in question.
2. Case data: a brief personality sketch so that the subject may perceive a similarity between Dick's described traits and the stereotype of a lawyer or engineer.

From the results, a consistent pattern emerged of subjects placing much greater evidential weight on the case data than on base-rate data. Judgments were typically resistant to changes in the ratio of lawyers to engineers in the population, and subjects typically ignored base-rate data even when the case data description was modified so as to become totally uninformative.

The explanation offered for this judgment bias by Tversky and Kahneman (1974, p. 1125) is that "subjects evaluated the likelihood that a particular description belonged to an engineer rather than to a lawyer by the degree to which this description was representative of the two stereotypes, with little or no regard for the prior probabilities of the categories."

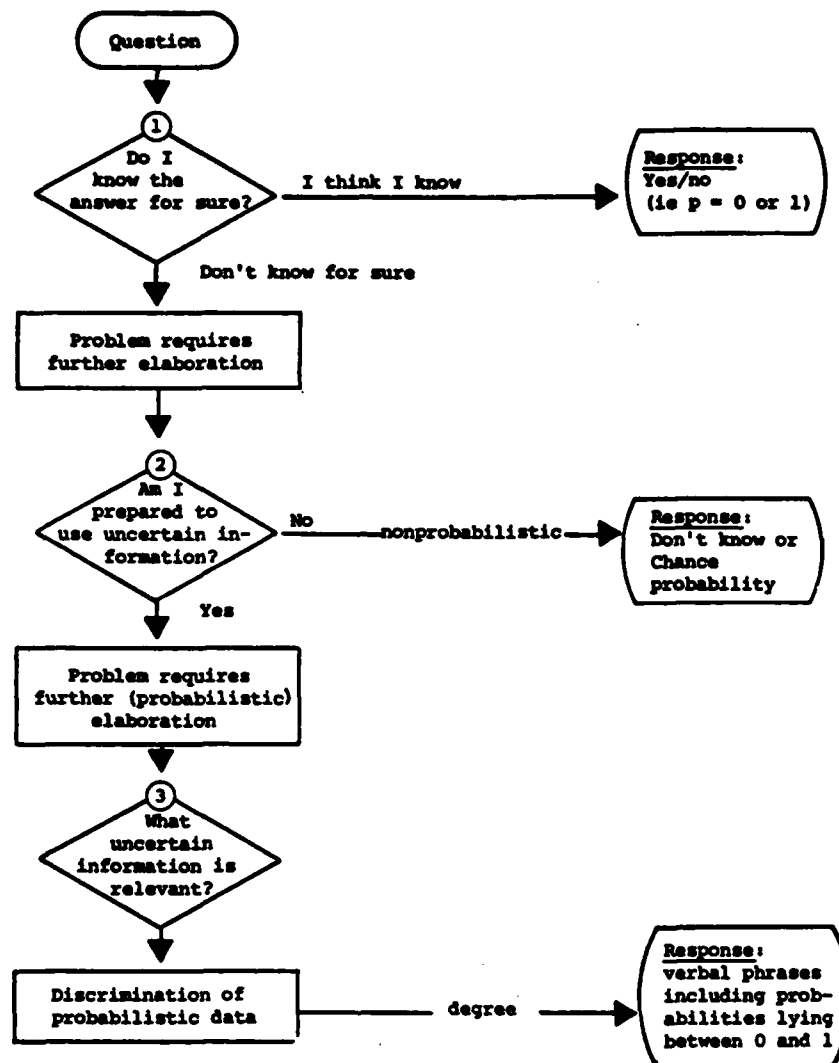
This explanation can be placed in context through the use of a model of sequential processes involved in probabilistic inference first presented in Phillips and Wright (1977). This model is reproduced as Figure 5. According to the Phillips and Wright model, the use of probabilistic information in a judgmental task requires a considerable elaboration of the problem over a simple deterministic assessment and involves the prior rejection of two other possible response modes as inappropriate: a response under certainty at stage 1 of the scheme shown in Figure 5, and a response consequent on a refusal to assess probabilistic evidence in spite of uncertainty at stage 2.

Because subjects in the Tversky and Kahneman experiments were presented with data inviting probabilistic inferences and were prepared, quite happily, to produce probabilistic responses, these subjects apparently elaborated the problem sufficiently to arrive at stage 3 of the Phillips and Wright schema. The biases typically found in the inferences made are thus the consequence of the particular way in which judges elaborate the structure of the problem at stage 3 in the schema. The judgmental problem is, in fact, modeled in a fashion that leads judges to be selective about which of the available probabilistic information they incorporate in the structure they use as a basis for their inferences.

Tversky and Kahneman's results indicate that subjects given case data perceived as being similar to (or representative of) a stereotyped class, at block 2 in Figure 6, typically made their judgment on the basis of the degree-of-fit between the case data and the characteristics of the stereotypes, neglecting base-rate data. Such judgments based on representativeness are represented in Figure 6 as those routes through the diagram marked by double lines. Thus the representativeness heuristic is characterized by two features:

1. Using a particular search instruction within semantic memory ("fit these data to the characteristics of a known class"), and
2. Making a probabilistic judgment on the basis of the degree of this fit without reference to any other available data.

It is the latter of these alone that violates the axiomatic model in failing to combine prior and posterior probabilities. Adherence to the axiomatic model would involve redrawing Figure 6 so that it became possible for a judge to move through both steps 2 and 3. In fact, judges given useful case data typically take the route passing through steps 1, 2, and 6 and so on (as a result of the representativeness heuristic). Judges given worthless case data



Notes.

- Stage 1. Judge's decision here is decided by a variety of factors classified under (a) cultural variables, (b) psychological variables, and (c) task variables.
- Stage 2. Judge's decision here is also decided by factors classified under cultural variables and psychological variables (especially those relating to open/closed mindedness; see Rokeach, 1960).
- Stage 3. This further decomposition of the basis for the inference involves the subjective elaboration of the structures represented in event and/or probability buffer systems shown in Figure 3.

Figure 5. Model of the cognitive processes involved in inference (developed from Wright & Phillips, 1977).


(data that permit little or no association with any already encoded class concept) persist in neglecting base-rate data and choose to take the route through steps 1, 2, and 4. Only subjects given no case data take account of frequencies, and they travel routes 1, 3, and 5.

It thus can be seen that judgments using case data and those using base-rate data involve traversing entirely different routes through the block diagram. A judgmental procedure involving the combination of case and base-rate data would require a different system from that of Figure 6, one that would open up routes to a problem solution quite different from those described previously.

We are now in a position to reformulate the problem of the representativeness effect and its etiology into the question, Why are judgments such as that described above typically structured through the sequential procedures of the form shown in Figure 6 rather than in a form permitting the combination of base-rate and case data? The following reformulation of the representativeness heuristic in production system terms suggests a solution to this question.

Applying the $R : C \rightarrow A$ formulation of a production system described in outline in section 4.1 (and in detail in Appendix A) to the structuring principles underlying Figure 6, we obtain:

R:	<u>Condition</u>	<u>Action</u>
	Active memory contains a class of objects that match this object or event.	Set the probability of this object or event being a member of the class in question to match the degree of fit between it and the matching class.



The implementation of this production leads the judge down routes 1, 2, and 6 and so on, in Figure 6, and thus entails the neglect of base-rate data and reliably reproduces the reported biases in responses (Tversky & Kahneman, 1974). In addition, the way in which memory is accessed in this production indicates that the site of the operation of the heuristic is the interface between memory and the probability buffering subsystem shown in Figure 3.

There is, however, a great deal more to the full specification of a production system than merely suggesting one possible production that such a system might contain. In particular, it is necessary for our purposes to understand the process whereby some productions pass a selection test and thus become potentially invocable, whereas others fail.

According to Anderson (1976, p. 186) production selection consists of "quick partial tests ... to see if the condition of a production is relevant to the current contents of active memory." A production can be implemented only if it is appropriate to the contents of active memory, and thus production selection is crucially dependent on the current activation of memory.

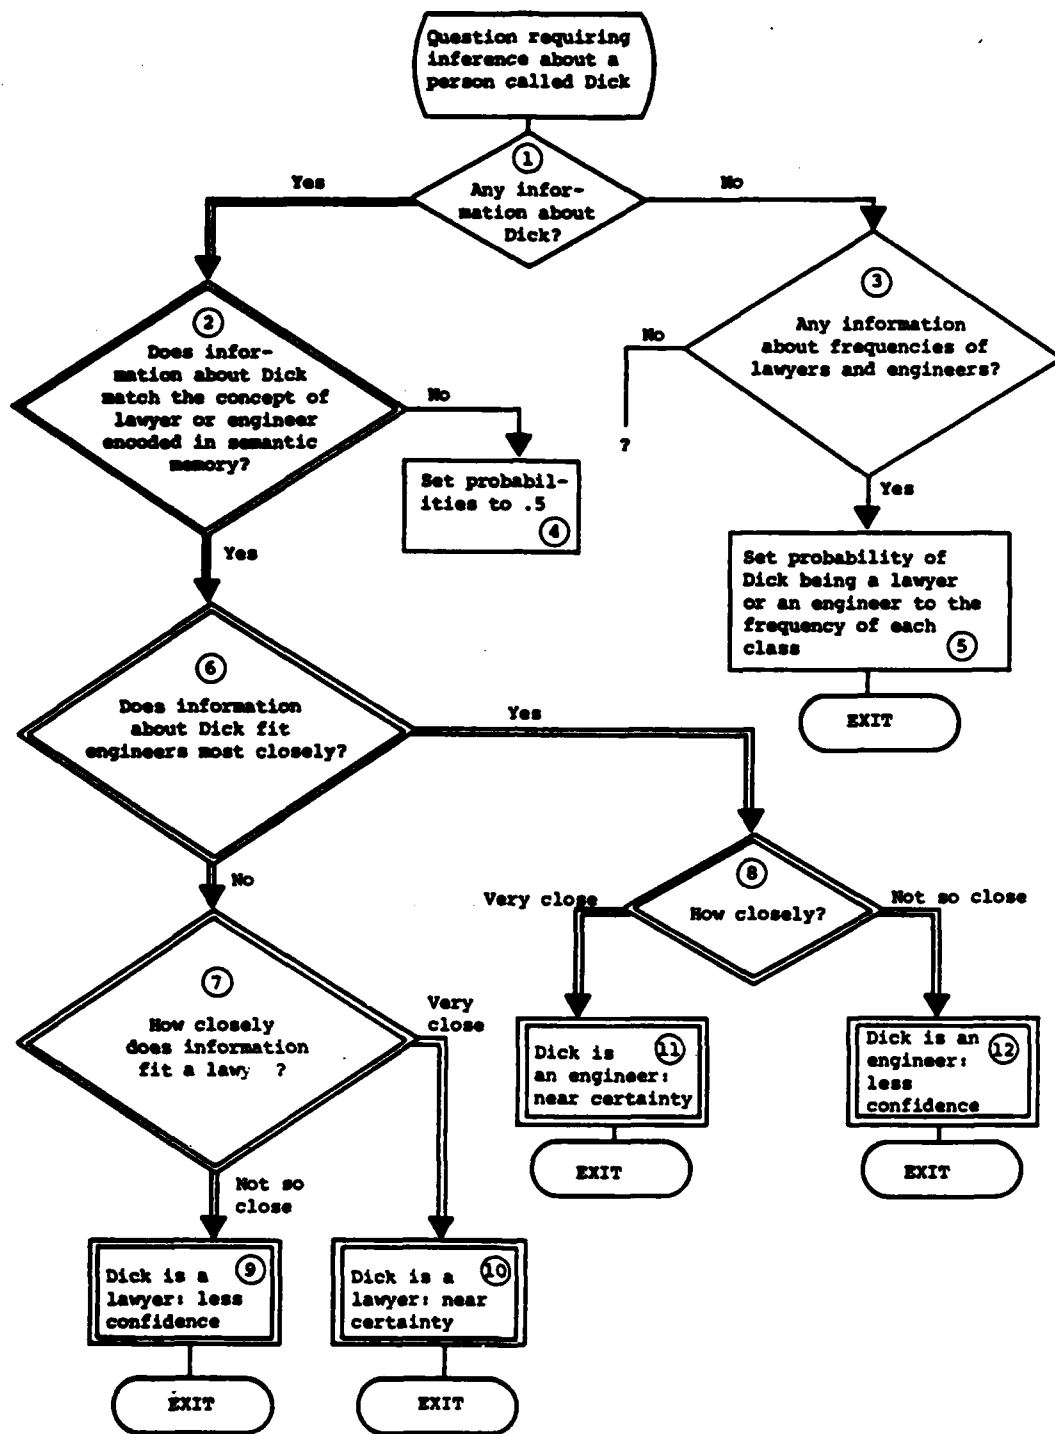


Figure 6. Block diagram of Tversky and Kahneman's (1974) model of judgment biased by representativeness.

We have so far shown only that if a particular production is implemented, then all the judgmental effects associated with representativeness will follow. It remains to be shown why this production is so commonly invoked. Tversky and Kahneman's work on representativeness contains no clues as to the determinants of its invocation. However, the crucial dependence of production selection upon the currently active memory indicates that the answer lies in discovering the determinants of the particular activation of memory. Newell and Simon (1972, p. 849) provide a clue to the answer:

We can know the objective task--"out there"--only through its particular representations. There is no neutral way of describing the task environment. As a consequence, task instructions do much more than define the task; they provide, in addition, a specific representation of it that can serve to define an initial problem space, and even parts of an initial problem solving program for the subject.

Thus, the repercussions of the subject's heuristic structuring of the task environment are felt throughout the entire judgmental process. One such repercussion, we suggest, is on the configuration of activated memory.

This accords neatly with our finding in the representativeness case. The process traced in Figure 6 was found to be the result of invoking the representativeness production. Any search to discover the basis for the invocation of that production leads back through the process of selection of potentially appropriate productions, and of memory activation, to the issue of the subject's heuristic structuring of the task environment. The whole process, from initial structuring to judgmental response, represents a highly constrained system, each stage of which is partially determined by prior stages.

What form of heuristic structuring precedes and determines a judgment by representativeness? While Tversky and Kahneman do not directly address the structuring issue, they have recently (1977) suggested a link between representativeness and the intuitive introduction of causal thinking into the judgmental process. This suggestion is, in effect, extended here to show how a subject's intuitive modeling of the task environment in causal terms determines which productions pass the selection test and consequently are available to be invoked. Which portions of memory a subject activates and in what form (in full, extended propositional form or in the more limited form required for matching-by-intersection tests) are determined by the subject's initial structuring of the task environment. Introducing a causal model of the task environment is one form of such initial structuring, and such a model can function as the initial memory activator. Abelson's (1976) work on cognitive scripts supports the belief in the predominance of this structuring of judgmental problems in causal terms.

Thus, in the example displayed in Figure 6, subjects' judgments of the probability of Dick being a lawyer typically involve the initial generation of a hypothesis concerning a causal relation between Dick's character traits and intentions on the one hand (derived from the personality sketch given) and Dick's choice of profession on the other, and this hypothesis creates an active partition of memory. It is assumed that the portion of memory likely to be activated in these circumstances, and hence available for accessing, consists of a simple pattern specifying traits and properties typical of "lawyer." The subject's judgmental task is now restricted to merely scanning

for a match between Dick's character traits derived from the given personality sketch and the properties associated with "lawyer" thus activated in memory.

Hence, the production typically implemented not only neglects base rates but engenders a fairly loose and imprecise assessment of case data. A more precise assessment of case data would require a propositional representation of relevant memory and a closer screening of the probe (that is, the character sketch of Dick) to make an exact assessment of the information it contains (for instance, encoding all the verbal elements in the proposition, as opposed to merely scanning for property words to be matched with properties encoded in memory). In short, the specific activation of memory involved in structuring the problem according to a causal model determines what information is accessible; this in turn determines what judgmental procedure is used to generate a probabilistic response. Methods based on Bayes' theorem provide a formal, optimal procedure for the combination of information in inference but are of necessity silent on how such information is accessed. However, in intuitive judgment, restrictions of access to information within memory (due to restricted active memory) can place crucial restraints on the inferential procedures that can be brought to bear. If a certain memory activation does not meet the condition of a production, that production is rendered inoperable.

The implication for decision analysis of the above interpretation of the representativeness heuristic is that the method of accessing information for input into the decision structure can prevent its optimal use in accordance with Bayes' theorem. If an inadequate or inappropriate structuring heuristic is employed at the interface between memory and the decision-making structure, the decision maker runs the risk of placing a structure on the decision problem that will of necessity commit him to inappropriate and nonoptimal inferences and actions, regardless of the optimality of the composition rules (Bayes' theorem, SEU, etc.), which he uses to manipulate information within the structure in arriving at his inference or plan of action.¹⁴

Incidentally, it is worth noting that judgment by representativeness counters *prima facie* the supposed ubiquity of simplification techniques designed to ease cognitive strain (Miller, 1956; Newell & Simon, 1972) in intuitive decision making. Since case data, unlike base-rate data, must be matched against information retrieved from memory to be put to use in the judgmental process, the diminution of cognitive strain thesis would find it surprising that base-rate data rather than case data are typically neglected. The former, it may seem, are more easily put to use, requiring less cognitive work on the part of the judge. The material just presented therefore suggests that it is misguided to explain the use of the representativeness heuristic in inferences as resulting from the need to reduce information processing load.

¹⁴ Sheppard (1976) and Humphreys (1979, section 5.2) present a case study of the use of such a heuristic by a division manager of a medium-large U.K.-based firm in making foreign location decisions and discuss the advantages and limitations of such an approach.

4.5 Classification of Heuristics

Table 1 summarizes and classifies the results of published investigations into the use of heuristics in intuitive judgment. Column 1 of the table specifies the features which, it is believed, are characteristic of intuitive judgment. Column 2 lists the explanations given for the presence of these features referenced by the authors supporting these explanations. In columns 1 and 2 we have used the currently prevalent nomenclature for the phenomena in question. In column 3, however, we have broken down this existing work on heuristics into two classes: those that may be relevant to the elicitation of structuring heuristics as we have defined the terms, and those that we believe are not relevant. We briefly state the basis for doing so in each case.

5. AN APPLICATION

Overview of MAUD

MAUD, Multiattribute Utility Decomposition, is an interactive, computer-based decision aid designed to help decision makers faced with a choice among alternatives where the basis for preference lies in differences in worth on a number of different attributes possessed in varying degrees by those alternatives. MAUD assists and guides the decision maker in (a) structuring and decomposition of such preferences in a multiattributed form and (b) finding out the tradeoffs he or she is prepared to make between values on the various attributes in recomposing these decomposed preferences into holistic utilities to be placed on the alternatives as a basis for choice. An embryonic version of MAUD was described by Humphreys and Humphreys (1975), and subsequent versions have been found to be of use in situations in which the decision maker has some intuitions about relevant aspects of the decision problem but has not, as yet, been able to discover its precise subjective worth structure.

MAUD was designed to work in direct interaction with the decision maker, without a decision analyst, counselor, or other expert as intermediary. However, since MAUD is limited to the examination of value tradeoffs among members of a homogeneous set of alternatives, a decision analyst or counselor in discussing a complex problem facing the decision maker should first arrive at an agreed definition of the set of alternatives whose worth structure MAUD is to investigate and the goal under which this worth structure is subsumed.

5.1 Example of the Use of a Heuristic Device: Use of Structuring Heuristics in the Elicitation of Poles of Attribute Dimensions

Decision theory is of necessity silent concerning the elicitation of attribute dimensions for incorporation in a decomposed preference structure. However, methods for eliciting such dimensions have been studied in some detail within research in the field of personality, stemming from the discussion by Kelly (1955) of the repertory grid, a device for conceptualizing an individual's dimensional cognitive structure. The rating form version of the repertory grid (see Bannister & Mair, 1968) is closely related structurally to the normal form decomposition of utilities of terminal events (called elements

Table 1

Classification of Heuristics

Uncertainty Structure:		
Reported effects in intuitive judgment	Explanation suggested by authors reporting the effect	Possible use as a component within structuring heuristics
(1) Conservatism in prior and posterior probabilities (DeCharme, 1970, claims conservative effects only at extreme ranges of posterior odds)	(1) Subjects misperceive data generator (Vlek & Beintema, 1967; Slovic & Lichtenstein, 1971) (2) Subjects misaggregate information (Edwards, 1968; Slovic & Lichtenstein, 1971) (3) Subjects tend to assign equal probabilities to all outcomes unless there is a reason to do otherwise (Pitz, 1975)	None--locates the effect at the interface between memory and environment None--locates effect within uncertainty structure None--locates effect within uncertainty structure
(2) Overconfidence (a) Neglect of base-rate data and failure to regress toward the mean in prediction from base rate and case data, unless base-rate data perceived as causally related to outcomes (Tversky & Kahneman, 1977) (b) Overconfidence in judgments of degree of fit between samples and populations (law of small numbers)	(1) Representativeness--judgment made by fitting case to a class of which the case is judged to be a representative member (Tversky & Kahneman, 1974; Kahneman & Tversky, 1972, 1973) (2) Subjects fail to measure predictive validity of data (Kahneman & Tversky, 1978) (3) Subjects misunderstand regressions and means (Peterson & Beach, 1967; Slovic & Lichtenstein, 1971) (1) Representativeness--subjects judge samples by the similarity to a model or prototype of population (Tversky & Kahneman, 1971, 1974; Kahneman & Tversky, 1972, 1973)	In cases in which subjects make inferences by referring case data to a parent population, representativeness is used as a search instruction for retrieval of information about parent population (see text) None--gives no indication of cognitive strategy used by the subject None--gives no indication of cognitive strategy used by the subject In cases in which a subject can retrieve a model of the population from memory, the probability of a sample matching the model is judged to be high, and hence confidence in the replicability of results is high

Table 1--Continued

Uncertainty Structure (continued):		
Reported effects in intuitive judgment	Explanation suggested by authors reporting the effect	Possible use as a component within structuring heuristics
	Subjects tend to explain consequent failures of prediction by causal accounts (Tversky & Kahneman, 1977)	Deviant results explained by searching "model of the world" for an intervention by a separate causal chain that causes the deviation from expectations.
(2)	In contexts characterized by lack of independent data, small samples can be highly representative of populations (Einhorn & Hogarth, 1978)	None--locates the effect in the environment
(c) Overconfidence in prediction from imagined scenarios	Availability--highly imaginable future developments are available for easy access but are not necessarily highly probable (Tversky & Kahneman, 1973, 1974)	If availability interpreted as determined by the structure of semantic memory, then--None--locates the effect within semantic memory If, however, availability interpreted as determined by the search instruction across the interface with semantic memory, then--effect is interpreted as a consequence of intuitive procedure for easy recall of information
(3) Judgment from correlations between data sets		
(a) Illusory correlation	(1) Subjects perceive causal relationship between data sets (Tversky & Kahneman, 1977)	Subjects introduce a causal model by which the data are explained and correlated
	(2) Subjects retrieval of information affected by availability (Tversky & Kahneman, 1974)	See entry in (2)(c), column 3

Table 1--Continued

Uncertainty Structure (continued):		
Reported effects in intuitive judgment	Explanation suggested by authors reporting the effect	Possible use as a component within structuring heuristics
(b) In correlation between input and output variables, judgment on basis of only 1 cell in a 2 x 2 data matrix	(1) Information from other cells often unavailable (Kishner & Hogarth, 1978) (2) Subjects lack an abstract notion of contingency (Svedlund, 1963)	None--locates the effect in the environment None--gives no indication of cognitive strategy used by the subject
(4) Greater redundancy of data --> greater confidence in inference from data (illusion of validity)	"Confidence is high when the data can be incorporated into a single integrated model that explains them"--causal schemata (Kahneman, 1974)	Data are associated with a causal model which, it is believed, causally determines the data
(5) The confidence that subjects have in their prediction depends primarily on the quality of match between the selected outcome and the input	Subjects select outcomes by the degree of their representativeness to the input (Tversky & Kahneman, 1974; Kahneman & Tversky, 1972, 1973)	This suggests an intuitive judgmental rule--select the outcome most representative of the input data, and a procedure for eliciting outcomes--access only outcomes representative of input data. The latter is of interest to us; the former is not (see Pitz, 1975, 1977)
(6) Inadequate adjustment of a probability estimate in the light of new information	Original judgment operates as an anchor, and adjustment from the anchor is typically insufficient (Slovic, 1972; Slovic, Fischhoff & Lichtenstein, 1976)	None--locates the effect within the uncertainty structure
(7) Relative invulnerability of opinions to conflicting evidence (the inertia effect)	Strong commitment to an hypothesis (Geller & Pitz, 1968; Slovic & Lichtenstein, 1971)	None--effect merely contrasts with Bayesian revision in uncertainty structure

Table 1--Continued

Uncertainty Structure (continued):		
Reported effects in intuitive judgment	Explanation suggested by authors reporting the effect	Possible use as a component within structuring heuristics
(8) Large deviations from mean of a probability distribution not weighted heavily (Peterson & Beach, 1967)	Subjects misperceive the impact of rare events (Slovic & Lichtenstein, 1971)	None--locates the effect within the uncertainty structure
(9) In assessment of predictive accuracy, subjects typically fail to learn from past failures of predictions	<p>(1) Hindsight bias--subjects tend to overestimate the accuracy of past predictions (Fischhoff & Beyth, 1975)</p> <p>(2) Nonindependence of outcomes due to treatment effects: poor judgment is vindicated by events (Zinbhorn & Hogarth, 1978)</p>	<p>None--on this account, subjects simply fail to use disconfirming evidence</p> <p>None--the effect is here considered as the result of lack of disconfirming feedback in the subject's environment</p>
(10) Greater confidence in prediction of variable B given variable A than vice versa, when A and B are equally informative	<p>Subjects hypothesize a causal relationship between A and B. The causal linkage $A \rightarrow B$ has greater evidential impact than the diagnostic relationship, $B \rightarrow A$ (Tversky & Kahneman, 1977)</p>	Information necessary for the construction of a causal model, which enable the intuitive assessment of uncertainties, must be retrieved from memory. The attempt to create a causal model may restrict the kind of information searched
(11) Subjects' recall of frequencies of events and sizes of populations typically biased by salience and recency	Availability--salience and recency of events and populations increase the probability of recall but do not increase probability of correct prediction (Tversky & Kahneman, 1973)	See entry in (2) (c), column 3.

Table 1--Continued

Utility Structure:		
Reported effects in intuitive judgments	Explanation suggested by authors reporting the effect	Possible use as a component within structuring heuristics
(1) Subjects prefer to compare outcomes on individual attributes serially rather than across attributes	Facilitates use of simplification techniques such as canceling out differences that are equal or nearly so, and reducing the number of attributes to be considered (Tversky, 1969; Slovic & Lichtenstein, 1971)	Gives some indication of conditions under which attributes are deleted from the utility structure
(2) Subjects tend to compare outcomes across attributes in more complex decision problems	This comparison procedure required when no common set of attributes between outcomes and when data about different outcomes not available simultaneously (Montgomery & Svenson, 1976)	Suggests that to access information about an outcome on all dimensions simultaneously, the decision maker must issue a fairly complex search instruction
(3) Representation of attribute dimensions of outcomes as bipolar scales (one end of the scale can be elicited by the method of triads: from three possible outcomes, choose word or phrase to describe an important respect in which two are similar and the third dissimilar) (Kelly, 1955)	<p>(1) Elicitation of the contrast end of bipolar scales by the difference method: give a word or phrase to describe how the third outcome is dissimilar from the other two (Kelly, 1955; Epting, Suchman, & Nickeson, 1971)</p> <p>(2) Elicitation of the contrast end of bipolar scales by the opposite method: give a word or phrase which states the opposite of the characteristic given for the likeness end (Kelly, 1955; Epting, Suchman, & Nickeson, 1971)</p>	<p>Gives a procedure for information search for structuring utilities</p> <p>Gives a procedure for information search for structuring utilities</p>

Table 1--Continued

Utility Structure (continued):		
Reported effects in intuitive judgments	Explanation suggested by authors reporting the effect	Possible use as a component within structuring heuristics
(4) The subject defines a subset of possible outcomes by making absolute evaluations on a single attribute that refer to a subjective criterion value	The elimination-by-aspects ^a (Tversky, 1972) rule used as a method for engineering a set of outcomes for which estimations of relative worths are then made, and which are subsequently linked, in the act-event structure, with specific immediate acts and uncertain events	Indicates how subjects access a manageable number of outcomes for consideration as consequences whose utility structures are to be traded off in the decision problem
(5) Increasing number of outcomes and attributes → decreasing proportion of aspects searched (Svenson, 1977)	Subjects presumably set criterion levels to restrict incorporation within utility structure to only the most salient or relevant outcomes and attributes of outcomes	None--does not help with the question of how salient outcomes and attributes are accessed
(6) Differences in weighting of dimensions, which are not explicable as differences in subjective importance of attributes	Attribute dimensions shared by all alternatives weighted more heavily than those shared by a subset (Slovic, 1972)	None--a feature of intuitive assessment within utility structure
(7) Subjects elicit possible outcomes or relevant attributes of outcomes by imaginative projection into possible futures	Scenario generation in utility structuring (O'Connor & Edwards, 1976; Janis & Mann, 1977)	Scenario generation may be a procedure not just for elicitation of outcomes but also for checking their realism, in the sense that an outcome may be deemed sufficiently realistic to merit consideration if it is derived from a well-defined scenario generation procedure (see, e.g., O'Connor & Edwards, 1976)

^a Tversky introduces EBA as a decision rule for choice of a preferred alternative from a set of possible options. Here EBA is used as a procedure for defining the set of possible options. While EBA used as a decision rule has been omitted from this list, we consider that this extended usage of EBA may be helpful for our purposes, and hence it has been included.

Table 1--Continued

Act-Event Structure:		
Reported effects in intuitive judgment	Explanation suggested by authors reporting the effect	Possible use as a component within structuring heuristics
(1) Subjects utilize only a subset of information available in the environment (Svenson, 1977)	(1) Subjects are using very simple decision rules (Payne, 1976) (2) Subjects decrease the number of choice alternatives and dimensions by the "method of successive limited comparisons" (Lindblom, 1964); they consider only those alternatives that differ slightly from present practices (Cyert & March, 1963)	None--locates the effect within the act-event structure A conservative heuristic, severely restricting information relevant for structuring decision tree
(2) Decision by external authority	Subjects choose between alternatives according to cultural or family tradition, or by prescription of an expert (Hogarth, 1974)	None---this simply notes reasons for failure to structure uncertainties in violation of axiomatic method
(3) Restructuring decision tree consequent on a reality check	Subjects divide attributes of the outcome of a choice alternative into two groups, A and B; from levels of A, predict levels of B; check predicted levels against observable levels; if a mismatch, delete the choice alternative (Humphreys, 1977)	This states a conservative rule specifying action to be taken in the event of a failure of a reality check and the conditions for applying the rule

Table 1--Continued

Act-Event Structure (continued):		
Reported effects in intuitive judgment	Explanation suggested by authors reporting the effect	Possible use as a component within structuring heuristics
(4) Because the subject sets a higher selling price for gamble A than gamble B —→ that the subject prefers A to B (Lichtenstein & Slovic, 1971)	Method of processing information is dependent upon the response the subject is required to make--e.g., setting a selling price or making a verbal choice. "Commensurability between a dimension and required response affects importance of that information in determining the response" (Slovic, 1972)	None--unless one speculates that response mode is a demand characteristic (Orne, 1962) of the task that determines the search instruction the subject uses to access information, a search instruction on which only information commensurable with the response mode is retrieved
(5) Decision affected by need to justify it to others	Justifiability--subjects take courses of action they believe will be readily explicable and justifiable to others (Slovic, Fischhoff, & Lichtenstein, 1976)	None--this decision procedure requires accessing information from memory but gives no indication how this access is achieved

within repertory grid terminology) into vectors of part-worth ratings on a set of attribute dimensions within a preference structure.¹⁵ The two principal heuristics discussed by Kelly (1955) for elicitation of attribute dimension poles are known as the Difference Method (DM) and the Opposite Method (OM). They are reviewed in Bannister and Mair (1968); Adams-Webber (1970); and Epting, Suchman, and Nickeson (1971). Both these methods are used in MAUD, an interactive computer program for the structuring, decomposition, and recomposition of preferences between multiattributed alternatives (Humphreys & Wisudha, 1979). When the decision maker has reached the point in structuring his or her decision problem at which at least three alternative terminal events are under consideration, and tradeoffs have to be made in determining relevant preference between these terminal events, MAUD may be used as an interactive aid in developing the preference structure among these alternatives. MAUD starts by using the Difference Method and a heuristic device to elicit an attribute dimension. Figure 7 reproduces the relevant portion of a printout from a typical MAUD interaction, together with comments on the functions of the various instructions used in achieving this elicitation.¹⁶

The material elicited through the use of this heuristic are two words (or phrases) that are assumed to define the poles of an attribute dimension scaled in terms of a monotonic part-worth preference function, so that ratings of the alternatives under consideration by the decision maker on this dimension will index the degree of part-worth¹⁷ of each alternative in terms of the degree to which they possess the relevant attribute.

Since the processes involved in the decision maker's semantic memory in forming the output shown in Figure 7 are not explicitly modeled, we cannot be sure that the elicited words have the required characteristics, as just outlined. However, we can perform a number of checks to test whether the necessary assumptions are met by the elicited words. These have been reviewed by Humphreys (1978a) as range of convenience (after Kelly, 1955), which establishes the possibility of actually making tradeoffs involving each and every alternative's rating on the attribute dimension; bipolarity, which requires that the poles be mutually exclusive and cumulatively exhaustive to insure an unambiguous unidimensional scale between the poles; and monotonicity, which insures that increasing numerical scale values index increasing preference.

¹⁵ These dimensions are called constructs within repertory grid terminology. However, the repertory grid decomposition does not represent a preference structure, because the ratings on constructs are J-scaled (linear between poles as defined, with no consideration of the position ideal point or the dimension) rather than I-scaled (part-worth, preference ratings). Coombs (1964), Humphreys (1977, section 4), and Humphreys and Wisudha (1979, section 3.5.1) discuss in detail the folding relationship between J-scaled and I-scaled ratings on an attribute dimension.

¹⁶ MAUD uses the Opposite Method in preference to the Difference Method when more than two attribute dimensions are already present in the preference structure. However, should the Opposite Method fail to elicit appropriate poles, MAUD will then (temporarily) revert to use of the Difference Method.

¹⁷ For a discussion of part-worth, see Kneppreth et al. (1974).

<u>Interaction</u> (user's responses are underlined)	<u>Comment</u>
<p>Can you specify a way in which one of the following:</p> <p style="margin-left: 40px;"> { 1 } LECTURER AT POLY { 2 } RESEARCH AT UNIVERSITY { 3 } MARKET RESEARCH </p> <p>differs from the other two (for your present purposes)? <u>YES</u></p> <p>What is the number beside the JOB that differs? <u>3</u></p> <p>You have said that MARKET RESEARCH is different from: LECTURER AT POLY and RESEARCH AT UNIVERSITY</p> <p>Please type one or more words on the same line describing a way in which MARKET RESEARCH differs:</p> <p>MARKET RESEARCH is: <u>NONACADEMIC</u></p> <p>On the other hand, LECTURER AT POLY and RESEARCH AT UNIVERSITY are: <u>ACADEMIC</u></p> <p>Are you reasonably happy with this description? <u>YES</u></p>	<p>Elicitation instruction designed to activate relevant portion of user's semantic memory five jobs were currently under consideration; MAUD chose three of them at random for use in this application of the Difference Method.</p> <p>Activation is successful.</p> <p>Confirmation of output from user's semantic memory by MAUD.</p> <p>Elicitation instruction requesting user's specification of poles of attribute dimension resulting from above activation of semantic memory (together with user's output).</p>
<p style="text-align: right;">↑ confirmation check</p>	

Figure 7. Interaction between decision maker and MAUD using the Difference Method to elicit poles of an attribute dimension.

If any of the checks of these assumptions fails, then the numbers assigned to alternatives on the offending attribute dimension will of necessity be incoherent within the utility buffer system shown in Figure 3 (and within any of the other decision-theoretic subsystems with which it is interfaced). To restore coherence, restructuring activity is required. Humphreys (1979) describes techniques designed to accomplish this as reordering techniques, as distinct from ordering techniques used to elicit initial (unchecked) components of the structure, similar to the Difference Method and Opposite Method outlined earlier.

Reordering techniques, like ordering techniques, involve crossing the interface with semantic memory and therefore involve heuristic devices. However, the various heuristics employed do not stand in isolation or in competition with one another; they can be conceived as integral parts of an ordering and reordering system that invokes heuristics as appropriate and checks the results, passing the results across the interface to the decision-theoretic subsystem if the checks are passed. If, on the other hand, one or more checks fail, other heuristics are invoked and a further set of checks made, with the system remaining active until a set of checks is passed in toto.

This type of system can be modeled as a production system. In the preceding section, we outlined a production system representation for heuristics involved in structuring operations. The next section shows how an embryonic form of such a system is currently employed in MAUD.¹⁸ MAUD's production system is probably the most advanced available in any current computer-based general purpose decision-aiding device, but it is still extremely limited.¹⁹ Before we can go further in building and testing production systems to handle structure ordering and reordering activities in decision aids involving transactions across the interface with semantic memory, we need a better understanding and specification of those heuristics that might be useful candidates for inclusion in such a system. Section 4 is concerned with the identification and classification of such heuristics.

5.2 Production System Control of Structure Ordering and Reordering Activities

This section outlines the way in which MAUD (Humphreys & Wisudha, 1979) uses a production system to control the structure ordering and reordering heuristics it employs in forming a preference structure described in terms of fully decomposed assessments of part-worths of a set of terminal events (outcomes) on utility-independent attribute dimensions. This system is, as yet, in its early infancy and has limited capability within a circumscribed area of the province of decision theory. However, it has proved very successful

¹⁸ See also Humphreys and Wisudha (1979, sections 2 and 3).

¹⁹ Other systems, such as OPINT (Selvidge, 1976; Allen et al., 1976) have comprehensive structure-reordering devices that operate within a decision-theoretic subsystem (the event buffer system in Figure 1 in the case of OPINT). Such systems operate according to normative specifications. While they are sometimes organized according to a production system, they are not considered here because they do not necessarily involve the invocation of heuristic devices.

in practice,²⁰ and this success has led us to propose that much more flexible, comprehensive systems be developed as soon as we have a better command of the building blocks for such systems: structuring heuristics of the type outlined in section 4.

Here we shall be concerned only with the control of the section of MAUD that deals with ordering and reordering the decomposed preference structure. It is interfaced with other sections of MAUD, which handle operations such as specification and revision of the set of terminal events among which preferences are to be traded off within the structure, elicitation, and revision of the relative value-wise importances of the attribute dimensions comprising the current preference structure and so on.²¹

Control is passed to this section of MAUD whenever a decision maker wishes to order (expand) or reorder (revise) the preference structure currently under consideration. (Even if this involves interrupting another task within the decision analysis, the ramifications of the interruption is handled at a higher level of control within MAUD.) Control within the section resides in an APPLYLIST²² of 15 productions. Each production (p) is of the form:

$$P : \{C \rightarrow A\}$$

where C is a vector of one or more conditions that must all obtain at the moment the APPLYLIST is scanned for that production to be implemented. A is a vector of one or more actions (procedures) that will be carried out by MAUD upon implementation of the production. Some, but not all, of these productions will involve interaction with the user. The actions transform the preference structure, and they also set appropriate condition flags during their execution. Certain condition flags may also be set at any time by the user overriding a request for particular input by MAUD with an interruption.²³

These productions, and their order of priority in the APPLYLIST, are shown in Figure 8. The conditions sensed and actions taken are defined as follows:²⁴

²⁰ See Humphreys (1978b) for a report of the use of this system with a wide range of decision makers and an analysis of gains made through its use.

²¹ See Humphreys and Wisudha (1979) for details of these operations.

²² APPLYLIST and other technical terms used in this section are defined in Appendix A.

²³ For instance, MAUD may request ratings of terminal events on an attribute dimension, but the user may reply that he or she is unhappy with the current definition of the dimension. The production system is designed to handle such interruptions in whatever way is most appropriate.

²⁴ The number in parentheses after each condition and action refers to the section in Humphreys and Wisudha (1979) in which these conditions and actions are described in detail.

<u>APPLYLIST</u>	<u>Comment</u>
ENTRY ↓	
P1 : {C6 → A7}	user-initiated deletion to
P2 : {C5 → A5}	part of preference structure
P3 : {C3 → A7, A3}	

P4 : {C8 → A13, A9, A10}	user-initiated change of content
P5 : {C4 → A9, A10}	within preference structure

P6 : {C15 → A16}	consequences of failure of MAUD-
P7 : {C14 → A15}	initiated check of adequacy of
P8 : {C13 → A9, A10}	decision-theoretic model

P9 : {C12 → A8, A11}	MAUD-initiated gathering of
P10 : {C11 → A12, A18, A14}	content within preference structure
P11 : {C10 → A5, A6, A17}	

P12 : {C2 → A4, A19}	user-initiated extension of
P13 : {C7 → A4, A19}	preference structure
P14 : {C1 & C2 & C16 → A2, A3}	
P15 : {C1 → A1, A3}	

↓	no further ordering-reordering
EXIT	operation required: task accomplished

Note. Productions are numbered in order of priority of execution within the APPLYLIST.

Figure 8. Productions used to control preference structure ordering and reordering activities within MAUD.

(1) Conditions

- C1 = The decision maker using MAUD (the user) wishes to have his or her semantic memory prodded to aid the elicitation of the preference structure relevant to the decision problem under consideration (i.e., tradeoffs between a set of terminal events, or outcomes leading to the formation of holistic preference values for those outcomes). (2)
- C2 = User is willing to define poles of an attribute dimension on the basis of the current (MAUD-generated) elicitation instruction. (2)
- C3 = User is not happy with current definition of poles of an attribute dimension. (2)
- C4 = User is not happy with current ratings on an attribute dimension. (2)
- C5 = User wishes to change current ratings on an attribute dimension. (2)
- C6 = User wishes to cancel/delete an attribute dimension from the current preference structure. (2)
- C7 = User wishes to add an attribute dimension to the current preference structure. (2)
- C8 = User is not happy with the current position of an ideal point on an attribute dimension. (3.5.1)
- C9 = More than two attributes are in current preference structure.
- C10 = Poles of an attribute dimension are defined, but ratings of terminal events (outcomes) on the scale spanning the poles are incomplete or absent. (3.2)
- C11 = J-scaled (raw) ratings of terminal events on a dimension exist but are not folded about the current ideal point on that dimension to give I-scaled (preference) ratings. (3.5.1)
- C12 = Ideal point is not currently set on an attribute dimension.
- C13 = Inadequate variance exists in I-scaled ratings on an attribute dimension. (3.5.2)
- C14 = Statistical independence check fails between I-scaled ratings on a pair of attribute dimensions. (3.3.2)
- C15 = Utility independence check fails in thought experiment conducted by MAUD in interaction with user. (3.3.2)
- C16 = Opposite Method heuristic for activating semantic memory failure flag not set.

(2) Actions (procedures executed by MAUD)

- A1 = Use Difference Method (DM) heuristic to activate semantic memory concerning attributes to be added to the current preference structure. (2)
- A2 = Use Opposite Method (OM) heuristic to activate semantic memory concerning attributes to be added to the current preference structure. (2)
- A3 = Investigate whether user wants to define poles of a new attribute dimension. (2)
- A4 = Elicit poles of new attribute dimension from user. (2)
- A5 = Elicit rating of terminal events (outcomes) currently under consideration as a J-scale spanning the poles of the current attribute dimension. (3.2)
- A6 = Investigate adequacy of J-scaled ratings on current attribute dimension. (3.2)
- A7 = Cancel/delete current attribute dimension. (2)
- A8 = Elicit ideal point on current attribute dimension. (3.5.1)
- A9 = Give user the option of canceling current attribute dimension. (2)
- A10 = Give user the option of changing his or her J-scaled ratings on current attribute dimension. (2)
- A11 = Check whether user is happy with the current position of his or her ideal point on a specified attribute dimension.
- A12 = Fold J-scaled ratings about the ideal point on an attribute dimension to form I-scaled ratings. (3.5.1)
- A13 = Cancel ideal point on specified attribute dimension.
- A14 = Check I-scaled ratings on current attribute dimension for statistical independence with I-scaled ratings on all other attribute dimensions active within current preference structure. (3.3.1)
- A15 = Conduct thought experiment with user to determine whether specified pair of attribute dimensions exhibit utility independence. (3.3.1)
- A16 = Delete specified (utility nonindependent) pair of attribute dimensions from current preference structures; elicit substitute pair of poles (defining new dimension within current preference structure). (3.3.1)

- A17 = Cancel Opposite Method heuristic failure flag.
- A18 = Check for adequate variance in I-scaled ratings on attribute dimensions. (3.3.2)
- A19 = Check whether user is happy with the definition of poles of current attribute dimension. (2)

- In operating the production system, the productions on APPLYLIST shown in Figure 8 are scanned in order of priority, always starting from the head of list. As soon as a condition specified in a production is matched to a condition currently obtaining, the action implied by that production is implemented in the sequence shown. On completion or interruption of the specified actions (either by MAUD or the user), control is immediately passed to the search of the APPLYLIST, which is scanned again. This cyclical procedure continues until the APPLYLIST is scanned completely with no production being activated. At this point, the structure ordering or reordering task is complete, and control is passed back to a higher level within MAUD.

This system has the great advantages of being flexible and capable of rapid expansion in future developments of MAUD, without running the risk of the control system getting out of hand or becoming indeterminate under particular conditions. It also gives us the ability to use and check heuristic procedures in an efficient way, let MAUD and the user share the direction of control, and still perform efficient housekeeping activities designed to minimize the extent of incoherence in the preference structure under development. Moreover, exit is not possible from the production system until the reordered preference structure is coherent; so the whole system can serve safely as a module for incorporation in a larger system, because it has only a single entry and exit and requires no external control.

The system is, of course, still limited, but our experience to date indicates that it will serve as a satisfactory basis for the development of much more comprehensive systems in the future, involving many more heuristic devices than those incorporated in the current version of MAUD.

6. STRUCTURING UNCERTAINTY

In section 3.1 we discussed two systems that can help the assessor determine a probability distribution for some target event--PIP systems and influence diagrams. Both systems were shown to act as buffers between the assessor's semantic memory and the core act-event subsystem. The systems facilitate probability assessment by providing structure, usually in disaggregated form, to uncertainty about the target event.

These are not the only two systems that can perform this function. As we reviewed systems for structuring uncertainty, it became clear that four fundamental strategies cover all the structuring systems currently in use, as well as those we could imagine might be helpful. We discuss these strategies in the next section, and then in section 6.2 we outline a new buffering subsystem.

6.1 Current Systems

An individual who must assess uncertainty about some target event can approach the problem in one of four ways, as shown in Figure 9.




Figure 9. Ways of assessing uncertainty about some target event.

Holistic judgment includes not only carefully considered judgment about the target event itself, but also judgment based on relevant historical data. For example, uncertainty about projected sales might be based on the past 12 months' sales figures, adjusted to take account of special conditions like inflation or new sales taxes that might prevail for the next 12 months. Strictly speaking, holistic judgments are not buffered; semantic memory is linked directly to the core act-event subsystem. However, holistic judgments are made frequently in most decision analysis, so this strategy for assessing uncertainty is included here for completeness.

When an assessor finds that uncertainty about the target event is difficult to assess because probabilities depend on other events, it may be useful to extend the conversation to include these other events. For example, if the probability assigned to event E is thought to be different if event F occurs than if it does not occur, then the assessor could be asked to assess the conditional probabilities $p(E|F)$ and $p(E|\bar{F})$, along with the unconditional probabilities $p(F)$ and $p(\bar{F})$. The probability associated with event E can then be calculated:

$$p(E) = p(E|F)p(F) + p(E|\bar{F})p(\bar{F})$$

If assessments of the probabilities associated with the occurrence and nonoccurrence of F prove difficult without considering event G, then a further extension of the conversation can be made to include event G. In this way, any number of related events can be considered. Technologies, or buffering subsystems, that are based on extending the conversation, include influence diagrams, event trees, and fault trees. At a theoretical level, there is no difference between these three technologies.

Two situations may lead the assessor to use Bayes' theorem. In extending the conversation, the assessor may be required to assess $p(F|G)$ but may find it easier to think about and assess the inverse probability $p(G|F)$. The earlier probability can be turned into the required probability by using Bayes' theorem, sometimes called the theorem of inverse probability:

$$p(F|G) = \frac{p(F)p(G|F)}{p(G)}$$

This type of structure is often referred to as "flipping the decision tree." The other situation in which Bayes' theorem is useful occurs when data are available that affect one's uncertainty about the target variable. The PIP system discussed in section 3.1 is an example; we might call this buffering subsystem a simple inference structure. Cascaded inference structures are also in use. Observable data reduce uncertainty about some unobservable indicator or factor, which in turn reduces uncertainty about the target variable. An example can be seen in insurance underwriting. A factory's trash disposal facilities can be directly observed. These, along with other data, give a fallible indication of the state of housekeeping, which is one of several factors that bear on the degree of fire risk posed by the factory. Another example is that of sensor uncertainty. A solid sonar return is a fallible indication of the presence of an enemy submarine, which in turn has some bearing on the extent of future hostilities between two adversaries.

Another way to disaggregate the target event is to determine the functional relationships the target event may have with other events. If the target event is an uncertain quantity, then it may be a function of other quantities, some of which are also uncertain. If the functional relationship between the target variable and these other variables can be determined, then assessments can be made of the uncertainty associated with the other variables and the probability distribution over the target variable calculated or determined by Monte Carlo simulation.

All of the uncertainty structuring systems we have encountered use one or more of these four strategies.²⁵

6.2 Buffering for Functional Relationships

When it is difficult to assess uncertainty about some target event because the event is complex, it may be possible to determine the functional relationship of the event to other, simpler events. Then uncertainty can more confidently be assessed for the simpler events. A probability distribution over the target event can be calculated if the functional relationships are simple enough for this procedure to be mathematically tractable or determined by Monte Carlo simulation.

What are the structural elements that could be used in this type of disaggregation? Typically, they are the same as those used in extending the conversation: events (and their outcomes), uncertain quantities and probabilities associated with the events, or uncertain quantities. This is the most elementary level of disaggregation familiar to decision analysis. However, an alternative level of disaggregation may often be more useful, especially when an interactive computer system is being used.

The approach can be illustrated with an example. One of the authors was asked to help the marine claims section of a large insurance company to improve its estimates of the eventual size of a settlement in cases of damage to the hull of a ship. Soon after an incident occurs, the claims department is notified of the accident. They are given only a brief description, but the information provides a rough basis for judging the eventual size of the settlement of the claim. There is often considerable uncertainty associated with this estimate. (The estimate is required for a variety of reasons; an example is that reserves must be set aside in a fund out of which claims are paid.)

After many hours of discussion with the head of the section, and following many revisions and alterations, the model shown in Figure 10 evolved. The eventual size of the settlement can be expressed as the product of four quantities: the net size of the claim, an image factor, a relationships factor, and a handling factor. Each of the last three factors acts as a multiplier that could increase the net size of the claim. For example, circumstances

²⁵ For example, Decisions and Designs, Inc., has produced TREE, an interactive computer program for decision tree modeling, and OPINT, an interactive program that includes a simple prestructured tree with a single target event whose uncertainty can be modeled with an influence diagram, supplemented (optionally) by a simple inference structure.

surrounding the claim may be adverse, which could lead the insurance company to pay out extra sums to preserve its image. Even if the circumstances are not adverse, the multiplier might be greater than one, though not so much as if the circumstances are adverse. Thus, uncertainty about the size of the factor is modeled in two stages: first, probabilities are assigned to the circumstances being adverse or not adverse; and second, two probability distributions over the factor are assessed, one conditional on the circumstances being adverse, the other on their not being adverse. Similar structures hold for the relationships and handling factors.

The net size of a claim is disaggregated into two quantities, the total cost less the net costs recoverable from a third party (if applicable). Total cost is further disaggregated into the size of the loss less a reduction for negligence plus the size of additional payments. Uncertainty about the size of the loss is expressed as a probability distribution over that quantity, which is bounded by the deductible at the lower end and by the amount of the indemnity at the upper end. A reduction for negligence may not apply; if it does, and this is assigned a probability, then the policy specifies that the size will be 10 percent of the size of the loss. Additional payments may not be necessary; if they are, for which another probability is assigned, then a probability distribution is assessed over the size of these payments.

Net costs recoverable from a third party are calculated by multiplying the total cost by a proportion of the total cost that is recoverable from a third party. Of course, costs may not be recoverable, but if they are (with some probability, which must be assessed), then a probability distribution over the proportion must be assessed.

In the figure, the shaded boxes identify the events and uncertain quantities for which probabilities must be assessed. Depending on the structure, assessments may be associated with discrete events, uncertain quantities, or uncertain quantities conditional on event outcomes.

Once the required assessments are made, a probability distribution over the eventual size of a settlement can be obtained by Monte Carlo simulation.

Note that each shaded box represents a structure based on the simple ingredients mentioned above: events and their outcomes, uncertain quantities, and probability distributions. However, some of the boxes are identical in structure: all the factor boxes are the same; the two bottom right boxes are also the same. Is it possible that, at this level of structuring, only a few structures are needed to represent most disaggregated target events? If so, then these few structures could be preprogrammed in a computer to serve as generic building blocks that would enable a user to model uncertainty about almost any target event.

6.3 Modular Uncertainty Structure

We believe that there are only a few basic structures; we call them modular uncertainty structures (MUSs). Although many more structures are, of course, possible, we think that those listed in Table 2 should be sufficient for most problems.

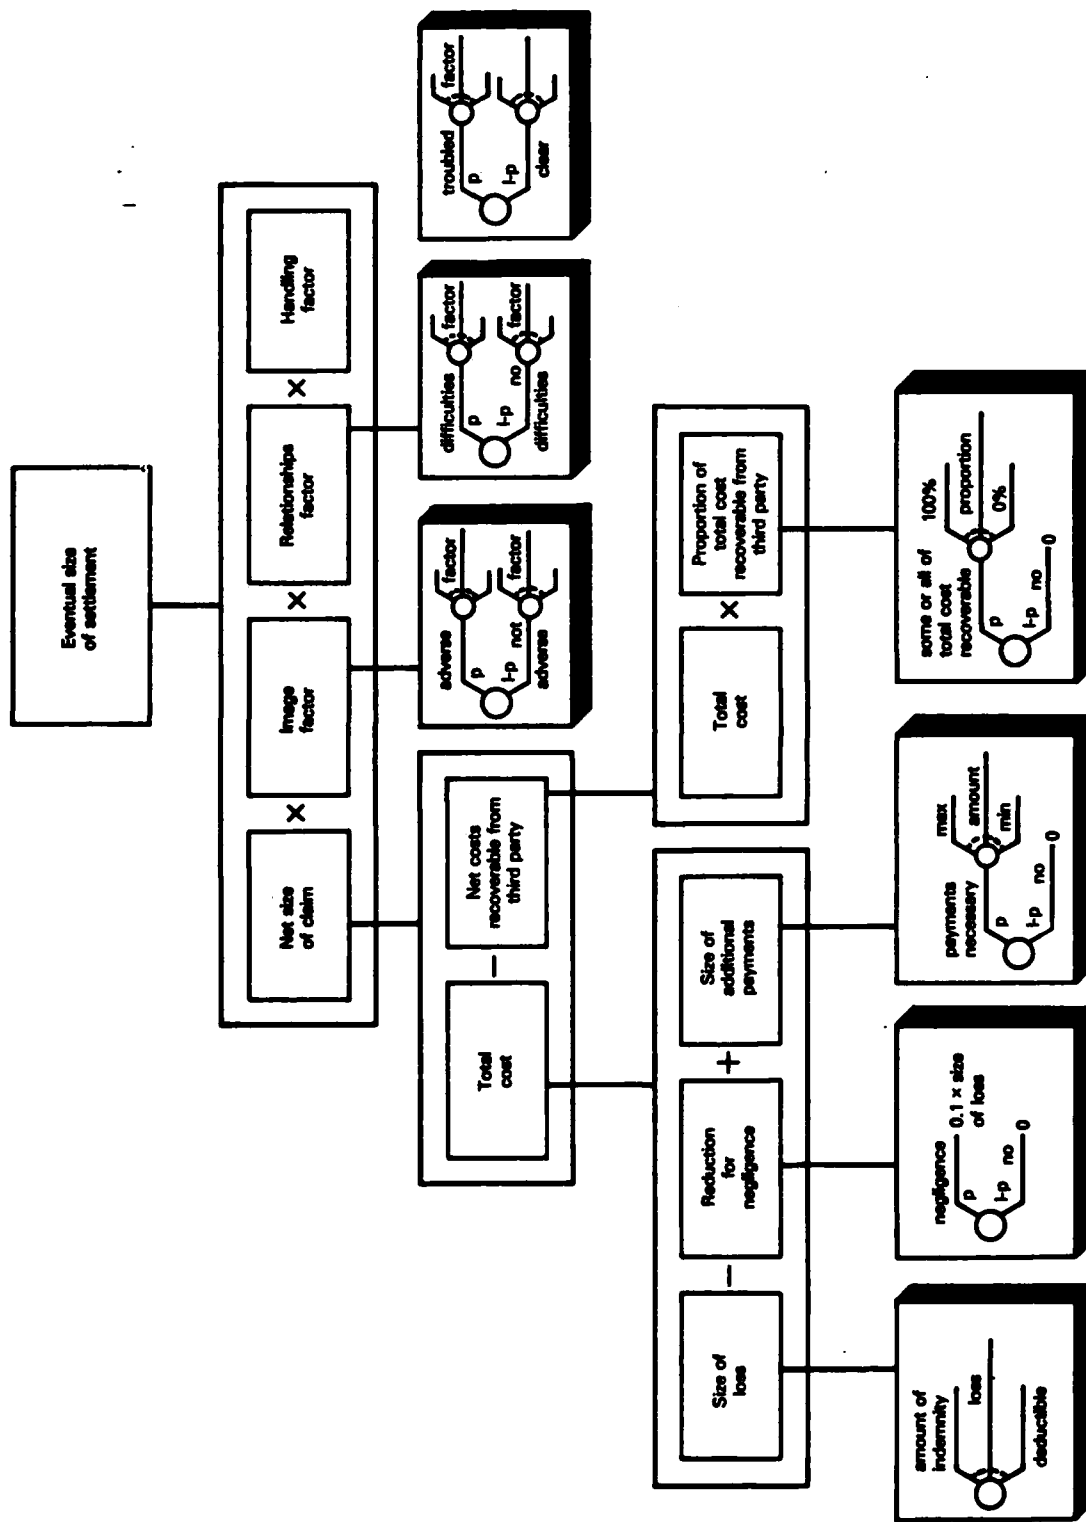


Figure 10. Modular uncertainty structure for hull claims model.


Table 2
Modular Uncertainty Structure

Inputs	Structure	Name
① $E, \bar{E}, C_E, C_{\bar{E}}, p$ either U_E and $U_{\bar{E}}$ or $C_E, C_{\bar{E}}$ and f		binomial distribution
② n, E_1, E_2, \dots, E_n p_1, p_2, \dots, p_n either U_1, U_2, \dots, U_n , or C_1, C_2, \dots, C_n and f .		multinomial distribution
③ selected fractiles, f, g (optional: max, min, truncation)		continuous distribution
④ selected fractiles, f, g , x, p . (optional as in 3)		mixed distribution
⑤ E, \bar{E}, C_E, p , fractiles, f, g		one conditional distribution

Inputs	Structure	Name
<p>⑥ E, E, p fractiles for x, f, g fractiles for y, h, i</p>		two conditional distributions
<p>⑦ n, E_1, E_2, \dots, E_n P_1, P_2, \dots, P_n fractiles for x, f, g fractiles for y, h, i \vdots fractiles for z, j, k</p>		n conditional distributions
<p>⑧ λ (rate parameter) C and f</p>		Poisson distribution with known consequences


Inputs	Structure	Name
--------	-----------	------

- ⑨ λ (rate parameter),
fractiles for C and f




poisson
distribution
with un-
certain
consequences

- ⑩ E, \bar{E}, F, \bar{F}
 $p(E), p(F|E), p(F|\bar{E})$





single
influence,
binary
events

- ⑪ Names of events and
their partitions,
conditional probabilities
for all linkages, uncon-
ditional probabilities
for lowest events.



multiple
influence

Inputs	Structure	Name
<p>12 Hypotheses, H; Data, D Prior probabilities, $p(H)$, likelihoods, $p(D H)$</p>		<p>simple inference</p>
<p>13 Hypotheses, H; Indicators, I; Data, D; prior probabilities, $p(H)$; $p(I H)$; $p(D I)$. (More precise speci- fication given in Kelly & Barclay, 1973)</p>		<p>hierarchical or cascaded inference</p>

Note. For cases 3-9, utilities may be expressed directly rather than as functions over consequences.

^aList of structural elements: E: event; p: probability; U: utility;
C: consequence; f,g,h,i,j,k: functions.

We can envisage a computer program that includes these 13 MUSs as the building blocks in a generic structure-building program. Any specific model could be built by defining disaggregated variables and their relationships, and then by assigning the appropriate MUS to each disaggregated variable. Each MUS used would then request the appropriate inputs from the user. Routines for assessing probability distributions could be included in the program to help the user generate coherent assessments.

It is possible that this program could be used by people who have some technical training but are not experts in decision analysis. By raising the structural level to the MUS from the rather molecular level used by decision analysts, buffering with semantic memory systems may be facilitated. The MUS structure may more closely represent internal structures built up by the expert in dealing with a particular class of problems over many years.

CONCLUSION

The reconceptualization of decision theory, presented in section 3 of this report, enabled us to outline the current limits of formal decision theory, and to go on in sections 4, 5, and 6 to explore beyond these limits, with a view to extending decision-theoretic methodology into the field of structuring decision problems. In section 4, we argued the case for the development of heuristic aids to structuring and suggested programmatic guidelines for this development; section 5 described the preliminary implementation of some such aiding devices within the structuring capability of MAUD. Section 6 pointed out a direction for further work in aiding the structuring of uncertainties by the use of modular uncertainty structures. The implications of the work described and the recommendations for future work are stated explicitly within each section.

What remains to be done is to contextualize these research efforts within an overview of the procedures involved in the process of decision structuring. We suggest such an overview next, decomposing the structuring process into a series of operations and suggesting where within this series the candidate structuring heuristics isolated in Table 1 of this report may be usefully employed.

Our review of the currently available work on structuring heuristics, summarized in Table 1, has led us to identify the following heuristics as worthy of further investigation as candidates for inclusion within structuring systems.

- Representativeness: a search instruction for making judgments about a present case by reference to a class encoded in memory.
- Causal schemata: A model of the world accessed from memory for judgment-making about present cases.
- Availability as a memory search instruction.
- Inter- and intra-attribute comparison and assessment of attributes of consequences.

- Search procedure determined by representation of attributes of consequences as bipolar scales.
- Elimination-by-aspects used as a procedure for reducing sets of options to be assessed.
- Systematic scenario generation: a strategy for eliciting realistic consequences of actions.
- Conservative heuristic for deleting actions from option set by comparing observed and predicted levels of attributes (Sheppard's investment manager's heuristic).
- A strategy dictating that only those alternatives approximating to present practice are included in the option set.

However, the specification of some candidates for inclusion within a decision-structuring system is merely one step toward the construction of such systems. What is needed in addition is a specification of a structuring system architecture into which each of these candidates may be fitted to aid the structuring of decision problems. In Figure 11 we present in general form a proposed architecture for such systems.

The roles of both the environment (represented as block 1 of Figure 11) and the constantly changing content of the decision maker's semantic memory (represented as block 2 of the figure) are explained in section 4 and in Appendix A. The right-hand column of blocks in Figure 11, presenting fairly gross descriptions of structuring operations necessary for providing the required inputs to a decision-theoretic evaluation model, concerns us here.

If the decision maker has recognized that a decision problem exists, we hypothesize that an interactive procedure incorporating the operations described in blocks 3 to 8 in Figure 11 will provide inputs crossing the interface with the decision-theoretic system adequate for assessment of a decision-theoretic model of the problem. Thus, the operation described in block 5 provides the option set to be evaluated within the decision-theoretic core system (see Figure 3). Block 6 structures those uncertainties arising within the decision problem necessary for numerical assessment within the decision-theoretic probability buffer system. Block 8 provides information concerning the criteria for evaluating consequences of actions necessary for the functioning of the decision-theoretic utility buffer system.

Blocks 3, 4, and 7, while not directly responsible for output into the decision-theoretic system, constitute the means whereby the decision maker constructs an internal representation of the problem requisite for the production of such output. The decision maker's construction of an understanding of his or her initial state (at block 3) is a necessary condition for the categorization of the decision problem at block 4, and it is by means of this categorization²⁶ that possible action plans for solving the problem are constructed at block 5. The decision maker's understanding of his goal, constructed at block 3, not only conditions the action plans, constructed at

²⁶ See von Winterfeldt, 1980.

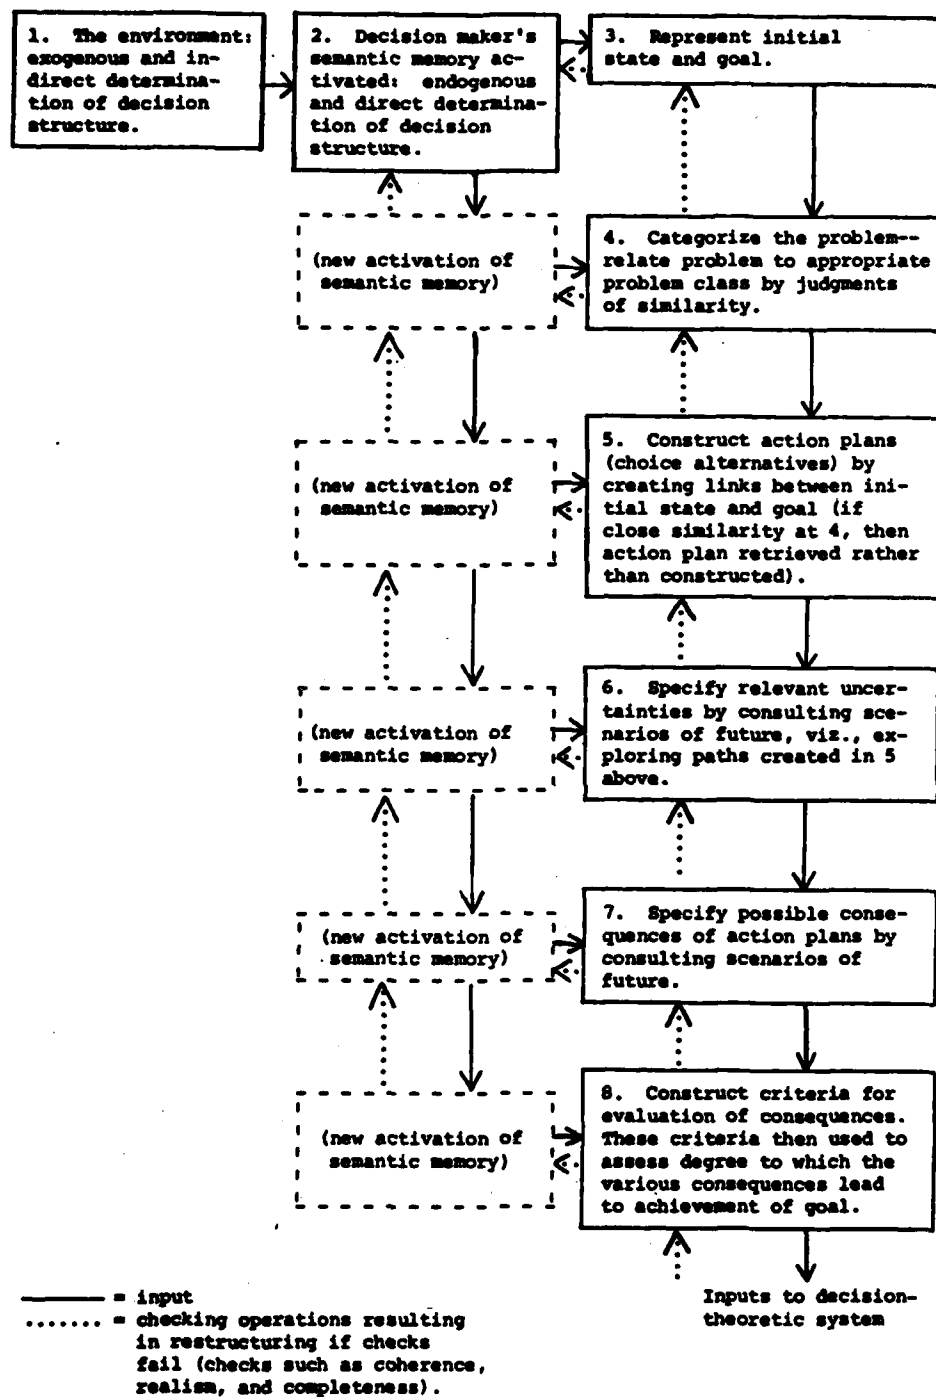


Figure 11. Diagram of basic operations required for creating a decision structure adequate for evaluation by the decision-theoretic system.

block 5, but also determines the criteria used for evaluation of the consequences of actions at block 8 in conjunction with the descriptions of consequences of actions at block 7.

If this constitutes a viable general description of a structuring system architecture, how might the candidate heuristics referred to above be employed within it? Each of blocks 3 to 8 in Figure 11 represents the site of operation of one or more of the candidate heuristics, as summarized in Table 3.

Rather than stating all of our reasons for assigning each heuristic to each structuring site in Table 3, we will concentrate on one example--that of representativeness. Representativeness may be a useful heuristic in circumstances in which it is appropriate to attempt to throw light on the present problem by reference to a class encoded in memory. It can thus be useful in categorizing a decision problem, represented as block 4 in Figure 11. However, in reinterpreting Tversky and Kahneman's research in representativeness in section 4.4, we have identified a further site of structuring operations at which representativeness may prove to be of help--namely in structuring the decision problem in such a way that uncertain events are assigned to classes, the subsequent numerical assessment of probability of the event being determined by the degree of fit between class and number. Thus it appears that representativeness also has a role in determining the output from block 6 of Figure 11.

These heuristics are in no way exhaustive of those we may need to consider in developing decision-structuring systems, and we consider that the best way to proceed is to start with systems of limited scope and expand the scope in interaction with decision makers facing a variety of decision-structuring problems. There is also a pressing need for a delineation of those structuring activities that could profitably be incorporated within a computer-based decision aid, and those that would reside better within the head of a decision analyst. It is our contention that aiding techniques that require knowledge of the world (or a simulation of the current contents of the decision maker's semantic memory) should not be automated because any device doing so would have also to be programmed with an enormous data base that would have to be constantly updated in the light of new information from every conceivable source. The human brain is likely to remain the best information processing and storage system of this type, at least within the foreseeable future.

We consider that such decision structuring systems should be controlled through a production system organization of the type now in use (in embryonic form) in MAUD and outlined in section 5.

Table 3

Sites of Operations of Structuring Heuristics

Site of structuring operations	Suggested heuristics operating at site
Represent initial state and goal (block 3)	Causal schemata Availability
Categorize the problem (block 4)	Representativeness Causal schemata Availability
Construct action plans (block 5)	Elimination-by-aspects Sheppard's investment manager's heuristic Access choice options approximating to present practice Availability
Specify relevant uncertainties (block 6)	Representativeness Causal schemata Systematic scenario generation Availability
Specify consequences of alternative actions (block 7)	Inter- and intra-attribute information search Systematic scenario generation Availability
Construct criteria for evaluation of consequences (block 8)	Search procedure determined by representation of criteria as bipolar scales

REFERENCES

- Abelson, R. P. Script processing in attitude formation and decision making. In J. S. Carroll & J. W. Payne (Eds.), Cognition and social behavior. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1976.
- Adams-Webber, J. R. Elicited versus provided constructs in repertory grid technique: A review. British Journal of Medical Psychology, 1970, 43, 349-354.
- Allen, J. J., Kelly, C. W., III, Phillips, L. D., & Stewart, R. R. Computer assisted option screening and intelligence assessments: Software user's guide. McLean, Va.: Decision & Designs, Inc., 1976.
- Anderson, J. R. Language, Memory and Thought. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1976.
- Anderson, J. R., & Bower, G.H. Human associative memory (2nd Ed.). New York: Hemisphere, 1974.
- Bannister, D., & Mair, J. M. M. The evaluation of personal constructs. London: Academic Press, 1968.
- Beach, L. R., Townes, B. D., Campbell, F. L., & Keating, G. W. Developing and testing an aid for birth planning decisions. Organizational Behavior and Human Performance, 1976, 15, 99-116.
- Bernoulli, D. Specimen theoriae novae de mensura, sortis. Commentarii Academiae Scientiarum Imperialis Petropolitanae, 1738, 5, 175-192. (Trans. by L. Sommer, Econometrica, 1954, 22, 23-36.)
- Brown, R. V. Heresy in decision analysis: Modeling subsequent acts without rollback. Decision Sciences, 1978, 9, 543-554.
- Brown, R. V., Kahr, A. S., & Peterson, C. R. Decision analysis for the manager. New York: Holt, Rinehart & Winston, 1974.
- Chinnis, J. O., Kelly, C. W., Minkler, R. D., & O'Connor, M. F. Single channel ground and airborne radio system (SINGARS) evaluation model. (Technical Report). McLean, Va.: Decisions & Designs, Inc., 1975.
- Cyert, R. M., & March, J. G. A behavioural theory of the firm. Englewood Cliffs, N.J.: Prentice-Hall, 1963.
- Davis, R., & King, J. An overview of production systems. Stanford, Calif.: Computer Science Department, Stanford University, 1975.
- Daves, R. M., & Corrigan, B. Linear models in decision making. Psychological Bulletin, 1974, 81, 95-106.
- DuCharme, W. M. A response bias explanation of conservative human inference. Journal of Experimental Psychology, 1970, 85, 66-74.

- DuCharme, W. M., & Peterson, C. R. Intuitive preference about normally distributed populations. Journal of Experimental Psychology, 1968, 78, 269-275.
- Edwards, W. Dynamic decision theory and probabilistic information processing. Human Factors, 1962, 4, 59-73.
- Edwards, W. Conservatism of human information processing. In B. Kleinmuntz (Ed.), Formal representation of human judgment. New York: Wiley, 1968.
- Edwards, W., Lindman, H., & Phillips, L. D. Emerging technologies for making decisions. In T. M. Newcomb (Ed.), New directions in psychology, II. New York: Holt, Rinehart & Winston, 1965.
- Edwards, W., & Phillips, L. D. Man as a transducer for probabilities in Bayesian command and control systems. In G. L. Bryan & M. W. Shelley (Eds.), Human Judgments and Optimality. New York: Wiley, 1964.
- Edwards, W., Phillips, L. D., Hays, W. L., & Goodman, B. C. Probabilistic information processing systems: Design and evaluation. IEEE Transactions on Systems Science and Cybernetics, 1968, SSC-4, 248-265.
- Einhorn, H. J., & Hogarth, R. M. Confidence in judgment: Persistence in the illusion of validity. Psychological Review, 1978, 5, 395-416.
- Epting, F. R., Suchman, D. I., & Nickeson, C. J. An evaluation of elicitation procedures for personal constructs. British Journal of Psychology, 1971, 62, 513-517.
- Fischer, G. W., Edwards, W., & Kelly, C. W. Decision theoretic aids for inference, evaluation and decision making: A review of research experience. (Technical Report). McLean, Va.: Decisions & Designs, Inc., 1978.
- Fischhoff, B., & Beyth, R. I knew it would happen: Remembered probabilities of once-future things. Organisational Behavior & Human Performance, 1975, 13, 1-16.
- Geller, E. S., & Pitz, G. F. Confidence and decision speed in revision of opinion. Organisational Behavior & Human Performance, 1968, 3, 190-201.
- Gustafson, D. H. Evaluation of probabilistic information processing in medical decision making. Organisational Behavior & Human Performance, 1969, 4, 20-34.
- Hammond, K. R., Kelley, K. J., Schneider, R. J., & Vancini, M. Clinical inference in nursing: Revising judgments. Nursing Research, 1967, 16, 36-45.
- Hogarth, R. M. Process training in clinical judgement. Behavioral Science, 1974, 19, 298-313.

- Howard, R. A. Decision analysis: Applied decision theory. In D. B. Hertz & Melese (Eds.), Proceedings of the Fourth International Conference on Operational Research. New York: Wiley-Interscience, 1966.
- Howell, W. C. Some principles for the design of decision systems: A review of six years of research on a command-control system simulation. Wright-Patterson Air Force Base, Ohio: Aerospace Medical Research Laboratories, Aerospace Medical Division, Air Force Systems Command, 1967 (AMRL-TR-67-136).
- Humphreys, P. C. Application of multiattribute utility theory. In H. Jungermann & G. de Zeeuw (Eds.), Decision making and change in human affairs. Amsterdam: Reidel, 1977.
- Humphreys, P. C. A review of some statistical properties of the repertory grid (and their cognitive implications). In Repertory grid methods. Leicester: British Psychological Society, 1978a.
- Humphreys, P. C. Multiattribute utility analysis of group decision making in "real" situations. (Final Report HR 3982/2.) London: Social Science Research Council, 1978b.
- Humphreys, P. C. Decision aids: Aiding decisions. In L. Sjöberg, T. Tyszka & J. A. Wise (Eds.), Decision analyses and decision processes, (Vol. 1). Lund: Doxa, 1979.
- Humphreys, P. C., & Humphreys, A. R. An investigation of subjective preference orderings for multiattributed alternatives. In D. Wendt & C. Viek (Eds.), Utility, probability and human decision making. Amsterdam: Reidel, 1975.
- Humphreys, P. C., & Wisudha, A. MAUD - An interactive computer program for the structuring, decomposition and recomposition of preferences between multiattributed alternatives. Uxbridge, Middlesex: Decision Analysis Unit, Brunel Institute of Organisation and Social Studies, Brunel University, 1979. (Technical Report 79-2.)
- Janis, I. L., & Mann, L. Decision making: A psychological analysis of conflict, choice and commitment. New York: Free Press, 1977.
- Kahneman, D. Cognitive limitations and public decision making. Paper presented to the International Symposium on Ethics in an Age of Pervasive Technology, 1974.
- Kahneman, D., & Tversky, A. Intuitive predictions: Biases and corrective procedures. ORI Bulletin, 1974, 12, 4.
- Kahneman, D., & Tversky, A. Subjective probability: A judgment of representativeness. Cognitive Psychology, 1972, 430-451.
- Kahneman, D., & Tversky, A. On the psychology of prediction. Psychological Review, 1973, 80, 237-251.

- Keeney, R. L. Energy policy and value trade-offs. (Research Memorandum RM-75-76.) Schloss Laxenburg: International Institute for Applied Systems Analysis, 1975.
- Kelly, C. W., & Barclay, S. A general Bayesian model for hierarchical inference. Organisational Behavior & Human Performance, 1973, 3, 388-404.
- Kelly, G. A. The psychology of personal constructs. New York: Norton, 1955.
- Kintsch, W. The representation of meaning in memory. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1974.
- Kneppreth, N. P., Gustafson, D. H., Leifer, R. P., & Johnson, E. M. Techniques for the assessment of worth. ARI Technical Paper 254, 1974. (AD78462a).
- Lee, W. Decision theory and human behavior. New York: Wiley, 1970.
- Lichtenstein, S., & Feeney, G. J. The importance of the data generating model in probability estimation. Organisational Behavior & Human Performance, 1968, 3, 62-67.
- Lichtenstein, S., & Slovic, P. Reversals of preferences between bids and choices in gambling decisions. Journal of Experimental Psychology, 1971, 89, 46-55.
- Lindblom, C. E. The science of muddling through. In W. J. Gore & J. W. Dyson (Eds.), The making of decisions. London: Collier-Macmillan, 1964.
- Luce, R. D., & Raiffa, H. Games and decisions: Introduction and critical survey. New York: Wiley, 1957.
- Miller, G. A. The magic number seven, plus or minus two: Some limits on our capacity for processing information. Psychological Review, 1956, 63, 81-97.
- Montgomery, H., & Svenson, O. On decision rules and information processing strategies for choices among multiattribute alternatives. Scandinavian Journal of Psychology, 1976, 17, 283-291.
- Newell, A. Production systems: Models of control structures. In W. G. Chase (Ed.), Visual information processing. New York: Academic Press, 1973.
- Newell, A., & Simon, H. A. Human problem solving. Englewood Cliffs, N.J.: Prentice-Hall, 1972.
- Norman, D. A., Rummelhart, D. E., and the CNR Research Group. Explanations in cognition. San Francisco: Freeman, 1975.

- O'Connor, M. P., & Edwards, W. On using scenarios in the evaluation of complex alternatives. (Technical Report 76-17.) McLean, Va.: Decisions & Designs, Inc., 1976.
- Orne, M. T. On the social psychology of psychological experiments: With particular reference to demand characteristics and their interpretations. American Psychologist, 1962, 776-783.
- Payne, J. W. Task complexity and contingent processing in decision making. Organisational Behavior and Human Performance, 1976, 16, 366-387.
- Peterson, C. R., & Beach, L. R. Man as an intuitive statistician. Psychological Bulletin, 1967, 68, 29-46.
- Peterson, C. R., DuCharme, W. M., & Edwards, W. Sampling distributions and probability revisions. Journal of Experimental Psychology, 1968, 76, 236-243.
- Peterson, C. R., Schneider, R. J., & Miller, A. J. Sample size and the revision of subjective probabilities. Journal of Experimental Psychology, 1965, 69, 522-527.
- Peterson, C. R., & Swensson, R. G. Intuitive statistical inferences about diffuse hypotheses. Organisational Behavior & Human Performance, 1968, 3, 1-11.
- Phillips, L. D., & Edwards, W. Conservatism in a simple probability inference task. Journal of Experimental Psychology, 1966, 72, 346-357.
- Phillips, L. D., & Wright, G. N. Cultural differences in viewing uncertainty and assessing probabilities. In H. Jungermann & G. de Zeeuw (Eds.), Decision making and change in human affairs. Amsterdam: Reidel, 1977.
- Pitz, G. F. A structural theory of uncertain knowledge. In D. Wendt & C. Vlek (Eds.), Utility, probability & human decision making. Amsterdam: Reidel, 1975.
- Pitz, G. F. Decision making and cognition. In H. Jungermann & G. de Zeeuw (Eds.), Decision making & change in human affairs. Amsterdam: Reidel, 1977.
- Pitz, G. F. Heuristic processes in decision making and judgment. Carbondale, Ill.: Southern Illinois University, 1979.
- Pitz, G. F., & Downing, L. Optimal behavior in a decision making task as a function of instructions and payoffs. Journal of Experimental Psychology, 1967, 21, 381-393.
- Post, E. L. Formal reductions of the general combinatorial decision problem. American Journal of Mathematics, 1943, 65, 197-268.
- Raiffa, H. Preferences for multiattributed alternatives. (Memorandum RM-5868-DOT/RC.) Santa Monica: The Rand Corporation, 1969.

- Rokeach, M. The open and closed mind. New York: Basic Books, 1960.
- Sandler, J., Dare, C., & Holden, A. The patient and the analyst. London: George Allen and Unwin, 1973.
- Savage, L. J. The foundations of statistics. New York: Wiley, 1954.
- Schum, D. A., & Kelly, C. W. A problem in cascaded inference: Determining the inferential impact of confirming and conflicting reports from several unreliable sources. Organisational Behavior & Human Performance, 1973, 3, 404-424.
- Selvidge, J. Rapid Screening of Decision Options. (Technical Report 76-12.) McLean, Va.: Decisions & Designs, Inc., 1976.
- Sheppard, J. M. Decision structuring under uncertainty: An empirical study of a foreign location decision. Working paper. London: London Graduate School of Business Studies, 1976.
- Slovic, P. From Shakespeare to Simon: speculations and some evidence about man's ability to process information. ORI Research Monograph, 1972, 12, 2.
- Slovic, P., Fischhoff, B., & Lichtenstein, S. Cognitive processes and societal risk taking. In J. S. Carroll & J. W. Payne (Eds.), Cognition and social behavior. Potomac: Lawrence Erlbaum Associates, 1976.
- Slovic, P., & Lichtenstein, S. Comparison of Bayesian and regression approaches to the study of information processing in judgement. Organisational Behavior & Human Performance, 1971, 6, 649-744.
- Slovic, P., & MacPhillamy, D. Dimensional commensurability and cue utilization in comparative judgement. ORI Research Bulletin, 1971, 11, 14.
- Smedslund, J. Concept of correlation in adults. Scandinavian Journal of Psychology, 1963, 4, 165-173.
- Stael von Holstein, C. S. A tutorial in decision analysis. In Readings in decision analysis. Menlo Park, Calif.: Decision Analysis Group, SRI International, 1977.
- Svenson, O. Process descriptions of decision making. In L. Sjöberg, T. Tyzka, & J. A. Wise (Eds.), Decision Analyses and Decision Processes (Vol. 2.). Lund: Doxa, 1979.
- Tversky, A. Intransitivity of preferences. Psychological Review, 1969, 76, 31-48.
- Tversky, A. Elimination by aspects: A theory of choice. Psychological Review, 1972, 79, 281-299.
- Tversky, A., & Kahneman, D. Belief in the law of small numbers. Psychological Bulletin, 1971, 76, 105-110.

- Tversky, A., & Kahneman, D. Availability: A heuristic for judging frequency and probability. Cognitive Psychology, 1973, 5, 207-232.
- Tversky, A., & Kahneman, D. Judgment under uncertainty: Heuristics and biases. Science, 1974, 185, 1124-1131.
- Tversky, A., & Kahnemann, D. Causal schemata in judgements under uncertainty. In M. Fishbein (Ed.), Progress in Social Psychology. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1977.
- Ulvila, J. W. A pilot survey of computer programs for decision analysis. (Technical Report 75-2.) McLean, Va.: Decisions & Designs, Inc.,
- Vlek, C. A. J., & Beintema, K. A. Subjective likelihoods in posterior probability estimation. (Psychological Institute Report No. E 014-67.) The Netherlands: University of Leiden, 1967.
- Wallsten, T. C. Processing information for decisions. In N. J. Castellan, Jr., D. Pisoni, and G. Potts (Eds.), Cognitive Theory (Vol. 2). Potomac, Md.: Lawrence Erlbaum Associates, 1978.
- Watson, S. R., & Brown, R. V. Case studies in the value of decision analysis. (Technical Report 75-10.) McLean, Va.: Decision & Designs, Inc., 1975.
- Wheeler, G., & Beach, L. R. Subjective sampling distributions and conservatism. Organisational Behavior and Human Performance, 1968, 3, 36-46.
- Winterfeldt, D. von. Structuring decision problems for decision analysis. Acta Psychologica, 1980, 45.
- Winterfeldt, D. von, & Fischer, G. W. Multiattribute utility theory: Models and assessment procedures. In D. Wendt & C. A. J. Vlek (Eds.), Utility, probability and human decision making. Amsterdam: Reidel, 1975.

APPENDIX A

PRODUCTION SYSTEMS AND SEMANTIC MEMORY

The production system (PS) provides a convenient tool for representation of methods, almost irrespective of how the researcher is disposed to analyze them. Whichever answers are favored to all the technical questions concerning the representation of knowledge, a specific PS system architecture can be designed for their specification. Whether a theory requires, for example, serial (Newell & Simon, 1972), or parallel (Anderson, 1976) implementing of processes, or a mixture of the two (Newell, 1973); whether its short-term memory data base is last-in first-out or first-in first-out; or whether an n-slot model of short-term memory (Newell, 1973) is used, or it is modeled as an active partition of an associative network in long-term memory (Anderson, 1976), the PS representation provides a precise and convenient modeling tool.

All the contents of any structured data base, of which that conventionally called long-term memory by experimental psychologists is of principal interest in personalist decision analysis, may be represented as procedural knowledge (knowing how) as in the Newell (1973) model. Thus Newell takes long-term memory to consist entirely of an ordered set of productions. Alternatively, our knowledge of procedures and our declarative knowledge of facts about the world may be separately represented. Anderson's (1976) model, which he calls ACT, is of this type. In ACT, declarative knowledge is represented as a propositional network encoded in semantic memory, while goal-seeking methods and procedures are specified as production systems. The propositional network is composed of complex configurations of nodes and links. Anderson (1976, p. 147) describes these as follows:

The nodes in the network such as ACT are intended to represent "ideas" The links represent access relationships or associations. That is, the links represent which ideas can lead to (elicit) each other.

Factual knowledge is represented as a set of propositions encoded in memory, the structure of each proposition being encoded as a set of nodes and relations (links) between nodes. Syntactic principles are invoked as a check upon the validity of propositions.

Heuristics, specified as production systems, operate on this propositional network, which constitutes the structured data base accessed by the decision maker. Production systems are invoked both to perform judgmental tasks (the output from the system) and to encode new propositions into the semantic memory network. The structure and content of the network both determine which productions can be invoked and provide the data structure that motivates the implementation of particular productions, in the manner outlined in Figure A-1.




Figure A-1. Steps in the operation of a production system processing inputs and outputs across interface between semantic memory and a decision-theoretic system. (The interfacing is shown in Figure 3.)

Productions are implemented as a result of a match between the conditions of productions and the contents of active memory. The matching check is performed on a subset of the total repertoire of productions available to the subject. This subset is composed of those productions that are potentially appropriate to the task environment and is labeled the APPLYLIST. An initial check of active memory is required to draw up the APPLYLIST. Hence, access to the propositional network is required at two separate stages of production implementation: in drawing up an APPLYLIST of potentially appropriate productions, and in the selection from the APPLYLIST of those productions to be implemented. However, access is restricted solely to that portion of the subject's memory that is currently active.

Any artificial model that purports to have psychological significance must restrict memory access so that all encoded elements are not accessible at any one moment in time. Newell (1973) achieves this by restricting the data structures that form the bases for the implementation of productions to a short-term memory. Anderson's (1976) model, on the other hand, is such that, while the implementation of a production system will draw further nodes and links into active memory, the selection of a particular production system for implementation is determined by an initial activation of memory, and it is this active portion alone that is accessible.

The structure and content of this activated portion of memory are the result of the decision maker's structuring of the task environment. The main report discusses how the activation is done, using detailed examples.